

PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION  
International Bureau

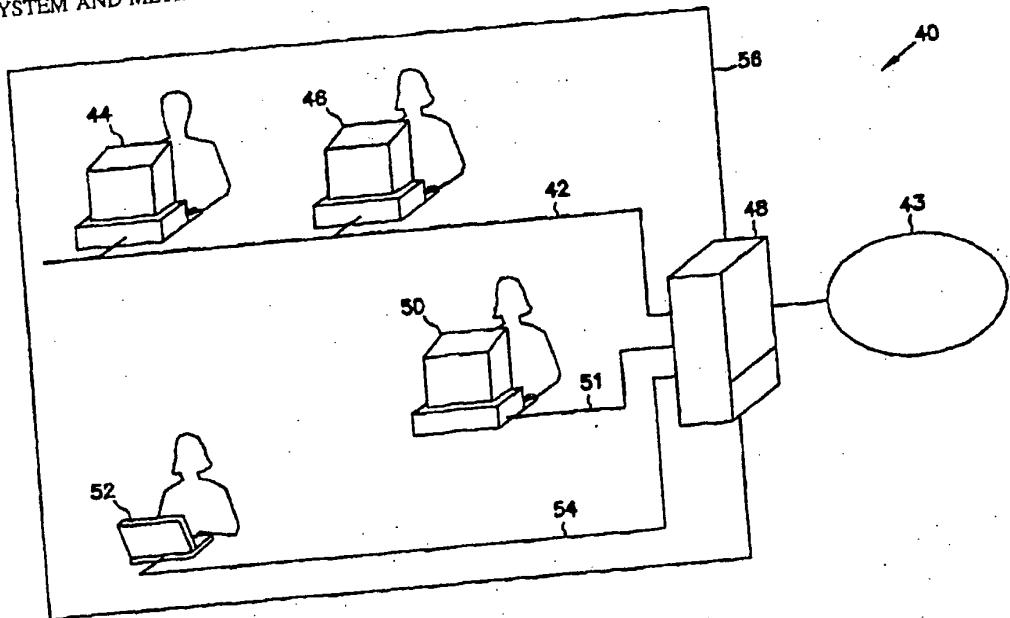


INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

WO 96/13113

(51) International Patent Classification 6 : H04L 29/06		A1	(11) International Publication Number: PCT/US95/12681	(43) International Publication Date: 12 October 1995 (12.10.95)	(11) International Publication Number: WO 96/13113	(43) International Publication Date: 2 May 1996 (02.05.96)
(21) International Application Number: PCT/US95/12681		(81) Designated States: AU, CA, JP, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).				
(22) International Filing Date: 12 October 1994 (12.10.94)		US				
(30) Priority Data: 08/322,078		Published With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.				
(71) Applicant: SECURE COMPUTING CORPORATION [US/US]; 2675 Long Lake Road, Roseville, MN 55113- 2536 (US).						
(72) Inventors: BOEBERT, William, E.; 4915 Dupont Avenue South, Minneapolis, MN 55409 (US). ROGERS, Clyde, O.; 5425 Township Drive, White Bear Lake, MN 55110 (US). ANDREAS, Glenn; 171 - 83rd Avenue N.E., Fridley, MN 55432 (US). HAMMOND, Scott, W.; 1292 Kohlman Avenue, Maplewood, MN 55109 (US). GOODERUM, Mark, P.; 3029 Blackstone Avenue South, St. Louis Park, MN 55416 (US).						
(74) Agent: VIKSNINS, Ann, S.; Schwegman, Lundberg & Woess- ner, 3500 IDS Center, 80 South Eighth Street, Minneapolis, MN 55402 (US).						

(54) Title: SYSTEM AND METHOD FOR PROVIDING SECURE INTERNETWORK SERVICES



(57) Abstract

A system and method for the secure transfer of data between a workstation (44, 46) connected to a private network (42) and a remote computer connected to an unsecured network (43). A secure computer (48) is inserted into the private network to serve as the gateway to the unsecured network and a client subsystem is added to the secure computer (48) in order to control the transfer of data from the workstation to the secure computer. The secure computer includes a private network interface connected to the private network, an unsecured network interface connected to the unsecured network, wherein the unsecured network interface includes means for encrypting data to be transferred from the first workstation to the remote computer and a server function for transferring data between the private network interface and the unsecured network interface.

BEST AVAILABLE COPY

**FOR THE PURPOSES OF INFORMATION ONLY**

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgyzstan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Larvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam

## SYSTEM AND METHOD FOR PROVIDING SECURE INTERNETWORK SERVICES

5

### Background of the Invention

#### Field of the Invention

The present invention relates to computer security, and more particularly, to an apparatus and method for providing secure access to a wide-area network.

10

#### Background Information

Advances in computer and communications technology have increased the free flow of information within networked computer systems. 15 While a boon to many, such a free flow of information can be disastrous to those systems which process sensitive or classified information. In a typical networked computer system, one or more workstations are connected over a network to a host computer or server. These workstations may range from low-cost personal computers to powerful UNIX processors. In such a system 20 the workstations, servers and even the connecting networks may all be at great risk of a security breach.

In developing a strategy for reducing the potential and consequences of a security breach (i.e. a computer security policy), one must assume that competent and dedicated individuals will mount active attacks on 25 the computer system's security mechanisms. These individuals are called the threat. The threat seeks to find vulnerabilities which can be exploited to cause a part of the computing system to operate in violation of its owner's security policy. Threats fall into two broad classes: Insiders and Outsiders.

Insiders are those individuals who have been granted some 30 level of legitimate privilege and then abuse that privilege. An example of an insider in the noncomputer world is a bookkeeper who uses his or her legitimate access to account records to embezzle. An example in the computer world is a systems administrator who uses his or her legitimate access to a computer system to generate fraudulent billings, payable to a corporation 35 owned by the administrator. Concern for insider actions also extends to

individuals who, through ignorance, incompetence or improper direction, cause security policy to be violated intentionally.

Outsiders are those individuals who have no legitimate privilege on the system but who can exploit vulnerabilities to gain access to it. An 5 example of an outsider in the noncomputer world is a burglar, who exploits weaknesses in locks and alarms to steal from a safe or lockbox. An example of an outsider in the network world is the "hacker" who takes control of a networked computer away from its legitimate owners.

The risk of security breach is compounded when a pathway is 10 provided from the internal, private network to an external wide-area network such as the Internet. The Internet is a loose conglomeration of networks connected by a standard network protocol. The lure of access to the Internet is the vast amounts of information that can be accessed by the user; the danger is that there are little or no controls on what individuals have access to 15 and what they may do with that access. Therefore, access to the Internet can provide an open door for exploitation of your own network by a variety of threats.

In effect, a wide-area network such as the Internet serves as a threat multiplier. Networks such as the Internet have evolved as fora for the 20 free exchange of ideas. This fact can be exploited by threats seeking to access or subvert a private network. For instance, the global connectivity of such a network means that data taken from a private network can be moved around the world very quickly. To compound this problem, the Internet contains a number of very large data archives which can be used to store data 25 transferred or posted from private networks. Hackers have also used the global connectivity of wide-area networks such as the Internet to directly manipulate computer facilities on the internal network (by such mechanisms as trying unlikely combinations of requests or commands) or to inject malicious software into the machine. Malicious software, which is able to do 30 the threat's bidding remotely and without direct control, can be injected manually or by such technical mechanisms as "viruses" or "worms." (One

such self-replicating piece of malicious software was responsible for a well publicized attack on computers connected to the Internet a few years ago.)

Internet protocols that have been developed to-date were not designed for security. For instance, Usenet news can be used by ignorant or 5 disgruntled employees to post company proprietary information in publicly accessible space. In some cases, this posting can be done anonymously (e.g. by using an anonymous file transfer mode or by posting the data to an anonymous server). In addition, the proprietary nature of data may be obscured by encrypting the data via one of a number of free, easily accessible 10 cryptographic packages.

In addition, since the standard Unix password is reusable, it is subject to capture and abuse by outsider threats. For instance, the use of 15 reusable passwords means that each password is vulnerable to being "sniffed out" and captured. Once captured the password can be used by an inside or an outside threat to gain access to a site. In addition, if the password belongs to someone with administrative privilege, the threat can use the captured password to gain administrative privileges on the internal network. The threat can then use that privilege to install a permanent "trapdoor" in order to ensure future access.

20 This combination of features makes the Internet particularly vulnerable to attack. A potential buyer of stolen information can anonymously post a solicitation along with his public key; potential sellers can then encipher the information desired with that public key and post it, secure in the knowledge that only the solicitor will be able to decipher it.

25 The existence of an active threat places requirements on a private network which are significantly different from the superficially similar problem of providing reliable service. A reliability engineer can take advantage of the low probability of certain phenomenon, and choose not to respond to them because they are so unlikely. A security engineer cannot do 30 this; a vulnerability, however obscure and unlikely, will be actively sought out by the threat, publicized to persons of like mind, and exploited over and over

once discovered. Countermeasures must therefore be developed which effectively close, or prevent the exploitation of, each system vulnerability.

A number of countermeasures have been proposed to reduce the vulnerability of networked systems. These countermeasures share three 5 characteristics:

1) *It takes a secret to keep a secret.* All information security mechanisms are based on the use of secrets which are shared by authorized individuals and kept from unauthorized ones. The secrets may be transformed, compressed or hidden inside protected hardware, but in every security 10 architecture there is one set of values, which, if known, would lead to the compromise of the whole system.

2) *Vulnerabilities always exist.* It is no more possible to achieve perfect security than it is to achieve perfect reliability; in fact, it is much less possible because you must assume that the threat is actively 15 working to discover the system vulnerabilities.

3) *Threats escalate continuously.* Installation of a given set of countermeasures does not eliminate the threat; it simply spurs it on to greater efforts to find ways of circumventing them.

These three common factors then pose the following problems 20 for the countermeasures engineer:

1) *Protecting the secrets that keep the secrets.* This is highest priority requirement, for loss of these values would lead to catastrophic breaches of security.

2) *Making vulnerabilities hard to find.* The embodiment of the 25 security mechanisms must be such that it is difficult for the threat to obtain details of their operation, or instances of them on which experiments may be performed.

The countermeasures proposed to date have focussed on either preventing the transfer of data or on encrypting the data using known 30 cryptographic methods in order to render it more difficult to compromise.

One method proposed for the prevention of unauthorized exploitation of the private network by inside or outside threats is an Internet

"firewall". "Firewalls" implement a security policy based on the routing information contained in individual packets transferred to and from the wide-area network. They look only at the headers of the packets and then make decisions based on where the packet is going and where it came from.

5      Typically, "firewalls" direct packets to a dedicated application machine which has a limited configuration of software. This application machine is then connected to a second router that limits its access to a specific set of internal systems.

10     A typical Internet "firewall" system 10 is shown in Fig.1. In Fig. 1, system 10 includes a router 12 connected over an internal network 14 to workstations 16 and 18. Router 12 is also connected to a wide-area network 20 such as the Internet. Router 12 runs Internet "firewall" software intended to inspect packet based traffic and remove or reroute packets meeting a predefined criteria.

15     "Firewalls" are header sensitive, not content sensitive. Therefore they are subject to various forms of attack. For instance, a hacker 22 may construct a packet having a header which looks like a header passed by the firewall. Such a packet will slip unnoticed past router 10 and onto one or more workstations 16, 18. In addition, a threat 24 may be able to access sensitive data on network 14 through the file transfer protocol ("FTP"). As noted above, a buyer 26 of stolen data may use Usenet news to solicit transfer of proprietary data from venal or disgruntled employees. Finally, a threat 28 may work in conjunction with a subverted employee 30 to transfer proprietary information via encrypted electronic mail or anonymous FTP.

20     Therefore, the Internet firewall approach has the following disadvantages:

- 1) This approach is vulnerable to attacks which construct fake header information (such as that by hacker 22 above). The theory of such attacks is well known; it is only a matter of time before turnkey scripts for mounting them become globally available on the Internet.
- 2) A "firewall" is an "all-or-nothing" approach to security. If an attacker gets through the "Firewall", then the internal network on the other

side lies naked and unprotected against effectively undetectable trojan horse attacks.

3) "Firewalls" can be difficult to configure correctly and even more difficult to keep secure because they have to be reconfigured as you 5 modify your internal network.

4) "Firewalls" cannot make security decisions based on data content, because they only see the data after it has been cut into packets and rearranged in the course of transmission.

5) "Firewalls" limit, in arbitrary and irrational ways, the user's 10 ability to interact with the Internet.

6) "Firewalls" require special "proxy" software for many Internet services. This means that there is a slow and costly development step required to "secure" a new service using the "Firewall" technique.

7) "Firewalls" require extra hardware and network connections, 15 which increases cost and administrative overhead.

The cryptographic countermeasures proposed to date have focussed on encrypting the data using known cryptographic methods in order to render it more difficult to compromise. Cryptography operates by performing mathematical transforms on data so that it is rendered 20 unintelligible to an outside observer. In order for the data to be retrieved, the transform is based on a second set of values called keying material. It is the keying material that is, in this case, the secret that keeps the secrets. Since both the writer and the authorized reader of the data must have equivalent 25 keying material, the central problem in cryptography is key management: the safe and reliable delivery of equivalent keying material to both ends of the writer-reader axis.

Cryptographic transforms use mathematical algorithms of great complexity and sophistication. In order to provide real-world security it is also necessary, however, that the embodiment or implementation of the 30 algorithm be not only correct but also free of vulnerabilities or side effects which can be exploited by the threat.

One commonly used class of cryptographic algorithms is called secret-key or symmetric. Such algorithms are called symmetric because the same element or value of keying material is used both to encipher (scramble) and to decipher (unscramble). They are called secret-key because that keying 5 material must be kept secret at both the writer and the reader ends of a communication. Secret-key systems require a some degree of prearrangement between the writer and the reader, so that the identical values of keying material are in place in advance of communication. As such, secret-key cryptography is most suited for communication amongst a closed community, 10 where membership in the community is known a priori. Simple changes in key distribution patterns can be used to add or delete individuals from the community.

Another class of cryptographic algorithms is called public-key or asymmetric. Such algorithms are called asymmetric because two 15 mathematically related elements of keying material are required: a public key, which is used to encipher but which cannot be used to decipher (unscramble), and a private key, which is the only value that can decipher. The corresponding private key, which is the secret that keeps the secret, is closely held. The public key, since it cannot be used to decipher, can be widely 20 disseminated. By this means a secret message can be sent without explicit prearrangement: the writer obtains the reader's public key from some service akin to a telephone directory, enciphers the message, and sends it with the knowledge that only the reader holds the private key that can decipher it.

A form of public-key algorithm can also be used to 25 authenticate, or sign, data. In this operation the private key is used to compute a value which is mathematically related to the data, called a digital signature. The private key is used so that only the holder of that private key can establish the distinctive value of the signature. The mathematics of the operation are such that the corresponding public can be used to determine the 30 validity of the signature. Thus only one person can sign, but any individual with access to the public key service can check the signature.

Public-key cryptography is most suited for communication within an open community, where it is desired to have secret and/or authenticated communication without prior arrangement. Adding individuals to the community is relatively simple, but deleting individuals is difficult.

5                   Cryptography has the following uses in information security:

- 1) Protection of communications links where the transmissions can be easily intercepted.
- 2) Protection of electronic mail where the messages may be forwarded through sites not under the control of the writer or the authorized 10 reader of the message.
- 3) Protection of data stored on removable media or media which is exposed to the possibility of physical theft.
- 4) Authentication, where the knowledge of a shared secret is used to verify the identity of an individual or a machine.

15                   The most sophisticated approaches to protecting data transferred over the unsecured Internet network are through the application of Global Cryptography at the Client workstation, so that data is enciphered at the source and deciphered at its destination. The principal application of this approach is to electronic mail. Global Cryptography can be implemented in 20 software, as in the Privacy Enhanced Mail system, or in personal tokens which combine the cryptographic mechanisms with an individual's certificate, as in the MOSAIC program.

25                   A less sophisticated approach is to apply the cryptography only on the wide-area network. Historically, there have been two ways to do this, called Link Encryption and End-to-End Encryption.

30                   In the Link Encryption approach, all bits coming out of a network node and onto the network are enciphered. This requires that the destination node have an identical cryptographic device and compatible keying material with the source. The disadvantage of link encryption is that all bits are encrypted, including those used to route packets over a packet-switched network. This effectively prevents a packet-switched network from working.

To permit the use of cryptography over packet-switched networks, the technique of End-to-End Encryption was devised. In this technique, only the packet contents are encrypted, and the critical routing information is left as plaintext. The "ends" in End-to-End encryption are 5 typically multi-user servers and not individual workstations, so that the problem of getting compatible keying material at each end is reduced to manageable proportions.

Neither data encryption nor the use of Internet "firewalls" address the array of vulnerabilities inherent to connection of an internal, 10 private network to an external, wide-area network such as the Internet. What is needed is a comprehensive and integrated security policy and apparatus for preventing exploitation of private network resources by both internal and external threats.

#### Summary of the Invention

15 The present invention provides a secure wide-area access system comprising a secure computer, an internal network and a workstation connected across the internal network to the secure computer. The secure computer comprises an internal network interface, a public network interface, public network program code used to communicate through the public 20 network interface to a public network, private network program code used to communicate through the internal network interface to the workstation and security policy program code for enforcing a Type Enforcement security mechanism to restrict access of a process to data.

According to another aspect of the present invention, a method 25 of protecting a computer system connected to an unsecured external network is described. The method comprises the steps of providing a secure computer, wherein the secure computer comprises security policy program code for enforcing a Type Enforcement security mechanism to restrict access of a process to data, connecting the Type Enforcement based secure computer to 30 the private network and establishing an assured pipeline for the transfer of data and programs between the private network and the external network through the secure computer. The step of establishing an assured pipeline

includes the steps of placing processes within domains, wherein the step of placing processes within domains includes the step of assigning processes received from the external network to an external domain, assigning types to files and restricting access by processes within the external domain to certain  
5 file types.

According to yet another aspect of the present invention, a secure server is described for use in controlling access to data stored within an internal network. The secure server comprises an administrative kernel and an operational kernel, wherein the operational kernel includes security policy  
10 program code for enforcing a Type Enforcement security mechanism to restrict access of a process received from the external network to data stored on the internal network and wherein the administrative kernel is restricted to execution only while isolated from the internal network.

According to yet another aspect of the present invention, the  
15 secure server comprises a processor, an internal network interface, connected to the processor, for communicating on an internal network and an external network interface, connected to the processor, for communicating on an external network. The processor includes server program code for transferring data between the internal and external network interfaces and security policy  
20 program code for enforcing a Type Enforcement security mechanism to restrict access of a process received from the external network to data stored on the internal network.

According to yet another aspect of the present invention, a system and method are described for the secure transfer of data between a  
25 workstation connected to a private network and a remote computer connected to an unsecured network. A secure computer is inserted into the private network to serve as the gateway to the unsecured network and a client subsystem is added to the workstation in order to control the transfer of data from the workstation to the secure computer. The secure computer includes a  
30 private network interface connected to the private network, an unsecured network interface connected to the unsecured network, wherein the unsecured network interface includes means for encrypting data to be transferred from

the first workstation to the remote computer and a server function for transferring data between the private network interface and the unsecured network interface.

According to yet another aspect of the present invention, a system is described for secure internetwork communication across an unsecured network. First and second secure computers are connected to first and second private networks, respectively, and to each other across the unsecured network. The first and second secure computers include a private network interface and an unsecured network interface for secure transfer of data from the first secure computer to the second secure computer over the unsecured network. The unsecured network interface includes means for encrypting data to be transferred from the first secure computer to the second secure computer. A client subsystem is added to workstations connected to each private network in order to control the transfer of data from the workstation to the respective secure computer.

#### Brief Description of the Drawings

Fig. 1 is a representation of a router-based "firewall";

Fig. 2 is a system level block diagram representation of a secure wide-area access system according to the present invention;

Fig. 3 is a more detailed block diagram representation of one embodiment of the networked computer system of Fig. 2;

Fig. 4 shows one embodiment of the system of Fig. 3;

Figs. 5a and 5b show the Type Enforcement mechanism used to prevent access, modification and/or execution of data objects without permission in a system such as that shown in Fig. 3;

Fig. 6 is a table of source code names of subtypes;

Fig. 7 is a table of access attributes;

Fig. 8 is a table of the new and effective Domains which result from particular syscalls;

Fig. 9 is a table listing the privileges which may be granted to a Domain;

Fig. 10 is a representation of steps taken in determining access privileges from the DDT;

Fig. 11 is a representation of steps taken in determining from the DIT the Domains a process can change to;

5 Fig. 12 is a system level block diagram of a wide area network connecting two organizational enclaves according to the present invention; and

Fig. 13 is a system level block diagram of another embodiment of a wide area network connecting two organizational enclaves according to the present invention.

10

#### Detailed Description of the Preferred Embodiments

In the following Detailed Description of the Preferred Embodiments, reference is made to the accompanying Drawings which form a part hereof, and in which are shown by way of illustration specific 15 embodiments in which the invention may be practiced. It is to be understood that other embodiments may be utilized and structural changes may be made without departing from the scope of the present invention.

A secure wide-area access system 40 is shown in Fig. 2. In Fig. 2, an internal network 42 connects workstations 44 and 46 to secure 20 computer 48. Internal network 42 is separated from a wide-area network 43 (such as the Internet) by secure computer 48. Secure computer 48 is also connected to a system administrator workstation 50 through a dedicated line 51 and to a workstation 52 through a serial interface 54. Secure computer 48 and workstations 44, 46, 50 and 52 make up an organizational enclave 56 of 25 data. The enclave is a "logical" enclave in that there is no requirement that the protected users and data be physically co-located, although such use of physical security measures is certainly possible.

It is important to isolate network 42 from network 43. To do this, secure computer 48 enforces an organizational security policy at the 30 interface between internal network 42 and wide-area network 43. It must do so in the face of active threat from both insiders and outsiders, whether by direct manipulation, the insertion of malicious software, or a combination of both. The system must protect its clients against attacks from wide-area

network 43, limit the damage done by subverted or incompetent clients, and be able to securely interact with clients of other systems 40 connected to wide-area network 43. It does this by surrounding the user with a set of protections that form organizational enclave 56.

Fig. 3 illustrates one embodiment of the secure wide-area access system 40 shown in Fig. 2. In Fig. 3, a workstation 63 (e.g. workstation 44, 46 or 52) connected to secure computer 48 over Private Network 64 (e.g. internal network 42 or serial interface 54) contains program code for communicating with secure computer 48 and through secure computer 48 to computers connected to wide-area network 43. Private Network 64 can be any means of communication, wired or wireless, which allows a workstation 63 to transfer data between the workstation and secure computer 48. In the example shown in Fig. 2, two embodiments of private network 64 are shown (internal network 42 and serial interface 54). It should be apparent that other embodiments of Private Network 64 can be implemented and the resulting system 40 would still fall within the scope of the present invention.

25 In one embodiment, the program code in workstation 63 includes a Client Interface Module 60 and a Client Protocol Module 62. Client Interface Module 60 accepts commands from, and displays results to, the user or Client. It can be embodied in a Graphical User Interface (GUI), a command line interface, or some combination of the two. Typical commands 30 would be to prepare an electronic message, examine incoming messages, request files from other sites, or any other operations typical of computer network usage.

Client Protocol Module 62 implements the protocol used to communicate between workstation 63 and secure computer 48. Client Protocol Module 62 can be implemented in either software or hardware, or a combination of both. In one embodiment, a Local Cryptography function 5 integrated into Protocol Module 62 has the specialized task of protecting and authenticating traffic on internal network 64 only. Different protocols and different cryptographic methods can be used for different Clients, depending on Client preferences and such factors as the nature of the physical connection (dialup, Local Area Network, etc.) between the Client Workstation and the 10 Secure Computer. It is most likely, though not required, that the closed nature of an organizational Client community (i.e. organizational enclave 56) will favor the use of secret-key cryptography in this module. In one embodiment, the Local Cryptography function is implemented in software in order to take advantage of software's flexibility and interoperability advantages over 15 hardware.

In another embodiment, the Local Cryptography function is implemented as a module separate from but operating in conjunction with Client Protocol Module 62.

In secure wide-area access system 40 of Fig. 3, program code 20 running on secure computer 48 is used to communicate through Private Network 64 to Client Protocol Module 62. In the embodiment shown in Fig. 3, the program code used to communicate with Client Protocol Module 62 is part of Private Network Protocol Module 66. In such an embodiment, Module 66 runs on secure computer 48 and interacts with Client Protocol Module 62 25 to provide protected and authenticated communication with workstation 63.

Likewise, program code running on secure computer 48 is used to communicate through a Public Network interface 72 to Public Network 74 (e.g. the Internet). In the embodiment shown in Fig. 3, the program code used to communicate with Public Network 74 is part of Public Network 30 Protocols and Cryptography Module 70. In such an embodiment, Module 70 runs on secure computer 48 and is used to provide protected and authenticated communication with individuals, sites, and other secure wide-area access

systems 40 on Public Network 74. Different protocols and cryptographic methods may be used when communicating with different entities on Public Network 74. It is most likely, though not required, that the open nature of Public Network 74 will favor the use of public-key cryptography in this 5 module.

Finally, program code running on secure computer 48 is used to implement servers and filter functions on secure computer 48. In the embodiment shown in Fig. 3, the program code used to implement the server and filter functions is part of Servers and Filters Countermeasures 68. As 10 such, the servers and filter countermeasures operate on the secure computer 48. They provide user services, such as the delivery of electronic mail or the transfer of data files and also enforce the organizational security policy by filtering the transfer of information and intercepting disallowed contents, labels, and/or addresses.

15

#### Cryptography in Secure Systems

The principal requirement for secure use of cryptography is a safe and reliable method for distribution of keying material. Reliability is as important as safety because if the material is not available then the users of 20 the system are faced with the unpleasant choice of either not using the cryptography (and thereby exposing their data to compromise or modification) or not transmitting. The key management requirements for a secret key system revolve around prearranged distribution of shared secrets. The key management requirements of public key systems revolve around insuring that 25 the writer of a document to be enciphered obtains the public key which corresponds to the reader's private key. Since the consequences of obtaining the wrong public key can be a breach of security, public keys are digitally signed by a notary or local authority who attests to their validity. Such signed public keys, with other optional information about the holder of the 30 corresponding private key, are called certificates.

Any effective key management system, and by extension any effective use of cryptography in a computer network, must also have facilities to solve the following problems:

- 1) Revocation. It must be possible to "take back" keying material so that an individual who was once authorized can have that authorization revoked.
- 5 2) Emergency rekey. It must be possible to "revive" the authorization of an individual if the keying material that grants the authorization is lost or destroyed.
- 10 3) Travelling user. The keying material that grants authorization to an individual must move around the network as the individual changes location.

Theoretically, the security of a cryptographic mechanism should rest only on the secrecy of critical keying material (all of it in a secret-key system, just the private part in a public-key system). As a practical matter, it is necessary to maintain protection of the mechanism for cryptography. This is especially true when the cryptographic device is partially or fully controlled by a computer system which may have been subverted through the use of malicious software. Such malicious software could cause the cryptographic 15 device to be bypassed either physically, by routing sensitive data around it, or logically, by causing a coherent pattern to be imposed on the timing or other characteristics of the output. This is not a cryptographic problem per se, but rather one that arises in the systems context of cryptography combined with potentially vulnerable computers.

20 25 In the secure wide-area access system 40 of Figs. 2 and 3, the burden of maintaining protection of the mechanism of cryptography is placed on secure computer 48. Secure computer 48 can be any type of machine whose features and/or implementation permits the operation of security-relevant functions to be trusted. Trusted computing systems have 30 been proposed for limiting access to classified information to those who have a sufficient level of clearance. Such systems depend on identifying the user, authenticating (through password, biometrics, etc.) the user's identity and

limiting that user's access to files to those files over which he or she has access rights. Such systems are described in U.S. Patent Nos. 4,621,321; 4,713,753; and 4,701,840 granted to Boebert et al. and assigned to the present assignee.

5           Typically, secure computers such as secure computer 48 provide safeguards through specialized hardware and software from direct attack on program code running in the secure computer. They have been developed to meet the following two objectives:

- 10           1) Limiting the privilege of users in a shared, or multiuser computer installation, so that malicious users cannot cause damage or compromise, and the effect of user error is minimized; and
- 15           2) Preventing damage or compromise that could result from the execution of malicious or erroneous software.

There have been two approaches to achieve the latter objective:  
15           exclusion, which seeks to prevent malicious software from entering the machine, and confinement, which allows the software into the machine and seeks to limit its effects. Existing secure computers fall into three broad classes:

- 20           1) Multilevel Secure Computers, which apply a confinement policy modelled on the U.S. Department of Defense system of data classification and personnel clearances. A Multi-Level Secure (MLS) Computer is capable of recognizing data of varying sensitivity and users of varying authorizations and ensuring that users gain access to only that data to which they are authorized. For example, an MLS computer can recognize the 25 difference between company proprietary and public data. It can also distinguish between users who are company employees and those who are customers. The MLS computer can therefore be used to ensure that company proprietary data is available only to users who are company employees.

- 30           2) Type Enforcing Secure Computers, which apply a confinement policy based on data flows through software subsystems in the machine.

3) Special Purpose Secure Computers, which apply an exclusion policy to insure that no malicious software is inserted in them, and then perform special-purpose security-related functions.

Secure wide-area access system 40 of Figs. 2 and 3 can make 5 use of any of these classes of machines, although it is most suited to being implemented on a Type Enforcing Secure Computer.

A freestanding Secure Computer has the following preconditions for secure use:

- 1) Protection of mechanism: the security mechanisms, 10 especially those embodied in software, must be protected from tampering or unauthorized modification. Since software mechanisms are prone to frequent update and improvement, there is a requirement for trusted distribution, that is, a means whereby administrators can be confident that the software they are installing is correct and proper.
- 15 2) User authentication: the security mechanisms often decide whether or not to allow an action based on the individual on whose behalf the action is being taken. There must be a method whereby the identity of a user can be authenticated.

In the case of a freestanding Secure Computer, physical 20 controls are typically sufficient to protect mechanism and simple methods such as passwords are sufficient to authenticate user identities. Designers of secure computers assume that unauthorized individuals will use a variety of means, such as malicious code and active and passive wiretaps, to circumvent its controls. Trusted subsystems of a secure computer must therefore be 25 designed to withstand malicious software executing on the untrusted subsystem, to confine the actions of malicious software and render it harmless. For instance, trusted computer systems based on host computers such as a Multilevel Secure (MLS) Computer make security breaches at the host computer more difficult by partitioning the system to isolate security 30 critical (trusted) subsystems from nonsecurity critical (untrusted) subsystems. In a similar manner, in Type Enforcing (TE) Secure Computers executables residing within the secure computer can only be executed if the person

requesting execution has execution privileges for that executable object. A further level of security can be achieved by preventing execution of any executable objects that have not been expressly recognized as a trusted executable by a trusted executable or by a system administrator.

5 In one embodiment of a TE-based system 40, only trusted executables are permitted to execute within secure computer 48. In such an embodiment, executables must first be reviewed and validated by a system administrator before they will be granted execution privileges on secure computer 48.

10 Secure computers do little, however, to prevent security breaches on the private network or at the workstation. One mechanism for avoiding such a breach is to authenticate the client to the secure computer over the network. The Local Cryptography function described above performs such a client authentication function. Another mechanism for avoiding a 15 network-related breach is to invoke a trusted path, a secure communications path between the user and the trusted subsystem. A properly designed trusted path ensures that information viewed or sent to the trusted subsystem is not copied or modified along the way. A trusted path authenticates not only the client to secure computer 48 (as in Local Cryptography above) but also 20 authenticates secure computer 48 to the client. As such, the trusted path mechanism guarantees that a communication path established between the trusted subsystem on secure computer 48 and the user cannot be emulated or listened to by malicious hardware or software.

Extension of the trusted path through the network to the user is, 25 however, difficult. As is described in a previously filed, commonly owned U.S. patent application entitled "Secure Computer Interface" (U.S. Patent No. 5,272,754 issued December 21, 1993 to William E. Boebert), "active" and "passive" network attacks can be used to breach network security. Active attacks are those in which masquerading "impostor" hardware or software is 30 inserted into the network communications link. For example, hardware might be inserted that emulates a user with extensive access privileges in order to access sensitive information. "Passive" network attacks include those in

which a device listens to data on the link, copies that data and sends it to another user. The '754 patent describes a system and method for ensuring secure data communications over an unsecured network. Operation of a trusted path in conjunction with an organizational enclave is described in U.S.

5 Patent No. 5,276,735, issued January 4, 1994 to Boebert et al.

In one embodiment, therefore, communication between Client Protocol Module 62 and Private Network Protocol Module 66 is made secure through the establishment of a Trusted Path between workstation 63 and secure computer 48 for all critical transfers.

10

#### Security Policy within the Secure Wide-Area Access System

The term security policy has acquired two meanings in the art:

15 1) The statement used by organizations and individuals to describe the objectives of their security activity, and to assign roles and responsibilities.

2) The rules used by a Secure Computer to determine whether or not certain actions may be performed.

In the latter case there are two kinds of policies:

20 2a) Label-based, in which the decisions are made on the basis of tag, or internal label, which is associated with a data object such as a file. The contents of the file are not examined by the decision-making mechanism.

2b) Content-based, in which the decisions are made on the basis of the contents of the file, message, or other data object.

Secure computers are required to perform the following tasks:

25 1) Protect data while it is being processed in unencrypted form. Certain operations, such as computations, editing, and transformation from one electronic message format to another can only be performed on data in unencrypted or cleartext form. Operations in encrypted, or ciphertext form, are generally limited to storage and transmission.

30 2) Enforce content-based security policies. Since such enforcement requires examination of contents, those contents must be in intelligible plaintext form.

3) Enforce individual roles and control the exercise of privilege. Cryptography inherently provides a binary or "all or nothing" privilege mechanism: either one possesses a decryption key, in which case one can read the data and then do whatever one pleases with it, or one does not possess the decryption key and operations on the data are prevented.

5

In a computer network, cryptography requires the following services from a Secure Computer:

10 1) Reliable and safe key management and distribution, including enforcement of limited roles for privileged individuals.

2) Protection of cryptographic mechanism from abuse by malicious software.

Correspondingly, Secure Computers require the following services from cryptography:

15 1) Authentication of user identities.

2) Protection of software mechanisms through trusted distribution.

3) Protection of data during storage or transmission in exposed environments such as a Public Network.

20 Underlying Principles of the Secure Wide-Area Access System

The first principle of system 40 is that the security services and alarms are centralized in a protected facility (secure computer 48) which is under the administrative control of a limited number of authorized individuals. Secure computer 48 can, and in general will, be physically protected to prevent unauthorized tampering or modification. In this way a greater degree of trust can be placed in its operation than in the operation of Client workstations 63, which are exposed, and in some cases portable.

25

Centralization means that security alarms are signalled only to authorized administrators who have privileges on secure computer 48; this facilitates response to insider attacks. Centralization also means that new services and countermeasures can be implemented simply by changing program code or

30

hardware on secure computer 48; such changes will be immediately available to, and imposed upon, all Clients on Private Network 64.

Secure wide-area access system 40 distinguishes between local authentication and protection, which takes place within the more protected 5 confines of a Private Network 64, and global authentication and protection, which takes place over a Public Network 74 shared with potentially hostile parties. All information is decrypted and examined in plaintext form by Filter Countermeasures 68 on secure computer 48. This permits the imposition of content-based organizational security policies and detailed audit of Client 10 interactions with Public Network 74. It also permits the intelligent transformation of data from one format to another when crossing the boundary between the Private Network 64 and Public Network 74. This ability is especially important in the case of electronic mail, where a large number of incompatible formats are in place.

15

#### A Type Enforcing Secure Wide-Area Access System

One embodiment of secure wide-area access system 40 of Fig. 3 is illustrated in the block diagram of Fig. 4. In Fig. 4, system 40 includes a secure computer 80 connected across a private network 82 to one or more 20 workstations 84. Workstations 84 are Intel-based IBM compatible personal computers running Windows 3.1 on the Microsoft DOS operating system. Protocol package 86 implements the protocol used to communicate between 25 workstation 84 and secure computer 80. In one embodiment, network 82 uses a TCP/IP protocol. In such an embodiment, protocol package 86 is a software package used to establish a WINSOCKET to network 82 on workstation 84. In one such embodiment, a Local Cryptography function is integrated into protocol package 88 in order to protect and authenticate traffic on network 82.

Client package 88 accepts commands from, and displays results to, the user or Client. It can be embodied in a Graphical User Interface (GUI), 30 a command line interface, or some combination of the two. Typical commands would be to prepare an electronic message, examine incoming

messages, request files from other sites, or any other operations typical of computer network usage.

In secure wide-area access system 40 of Fig. 4, program code running on secure computer 80 is used to communicate through Private Network 82 to protocol package 86. In one embodiment, secure computer 80 is an Intel Pentium-based machine running a hardened form of BSD386 Unix. A system based on a 90 MHz pentium with 32 megabytes of memory, 2 gigabytes of hard disk space, a DAT tape for backup and a CD-ROM for software loads has been found to be adequate.

10 In the embodiment shown in Fig. 4, the program code used to communicate with protocol package 86 is part of protocol package 90. In such an embodiment, package 90 runs on secure computer 80 and interacts with protocol package 86 to provide protected and authenticated communication with workstation 84. For instance, a Local Cryptography 15 function may consist of software which executes on workstation 84 to establish client authentication at login. In such a system, when a user logs into network 82, a message is sent from workstation 84 to secure computer 80. Secure computer 80 responds with a number (in one embodiment, this is a seven digit number) which is sent unencrypted to protocol package 86 on 20 workstation 84. Protocol package 86 then generates a request, through client package 88, to the user to enter his or her personal identification number (PIN). Protocol package 86 takes the PIN and combines it with a predefined number stored on workstation 84 to form a DES encryption key. That DES 25 encryption key is then used to encrypt the number received from secure computer. The encrypted number is sent to secure computer 80, where it is decrypted. If the correct machine number and PIN number were used for that particular user, secure computer 80 will be able to reconstruct exactly the number it sent to workstation 84. If not, an error is generated and an entry is made in the audit log. In one embodiment, active spoofing countermeasures 30 are then executed in an attempt to keep the threat in the vicinity of workstation 84.

Once the client is authenticated, communication on network 82 is in clear text.

Likewise, program code running on secure computer 80 is used to communicate through a Public Network interface to a Public Network. In 5 the example shown in Fig. 4, the public network is the Internet. In such an embodiment, the program code used to communicate with the Internet is part of an Internet protocols 94 which communicates with computers on the Internet through Internet connection 96. Internet protocols 94 runs on secure computer 80 and is used to provide protected and authenticated 10 communication with individuals, sites, and other secure wide-area access systems 40 over the Internet. Different protocols and cryptographic methods may be used when communicating with different entities on the Internet. In one embodiment, a tcp wrapper package operating in Internet protocols 94 is used to sit on the external, public network so that information about external 15 probes can be logged. It is most likely that the open nature of Public Network 74 will favor the use of public-key cryptography in this module.

Finally, program code running on secure computer 80 is used to implement servers and filter functions on secure computer 80. In the example shown in Fig. 4, the program code used to implement the server and filter 20 functions is part of Internet Servers and Filters 92. As such, the servers and filter countermeasures operate on secure computer 80. They provide user services, such as the delivery of electronic mail or the transfer of data files and also enforce the organizational security policy by filtering the transfer of information and intercepting disallowed contents, labels, and/or addresses.

25 As noted above, in one embodiment secure computer 80 is an Intel Pentium-based machine running a hardened form of Berkeley's BSD386 Unix. In that embodiment, BSD386 is hardened by adding a Type Enforcement mechanism which restricts the access of processes to data. Type Enforcement operates in conjunction with page access control bits in the 30 virtual page translator of the Pentium to control access to objects stored in secure computer 80 memory. To accomplish this, system calls in the basic BSD386 kernel were modified as shown later in this document so that Type

Enforcement checks cannot be avoided. Certain other system calls were either disabled or had certain options disabled.

In the hardened BSD386 according to the present invention, Type Enforcement controls are enforced by the kernel and cannot be circumvented by applications. Type Enforcement is used to implement data flow structures called *Assured Pipelines*. Assured pipelines are made possible by the so-called "small process" model of computation used by Unix. In this model, a computational task is divided up into small virtual units that run in parallel to each other. Unix provides a crude and loosely-controlled way of sharing data between processes. Type Enforcement supplants this with the rigorously controlled, configurable structure of assured pipelines.

In addition, secure computer 80 has been configured under BSD386 to run in one of two states: administrative and operational. In the administrative state all network connections are disabled and the Server will only accept commands from a properly authenticated System Administrator accessing the system from the hard-wired administrative terminal (such as terminal or workstation 50 in Fig. 2). This feature prevents anyone other than the System Administrator from altering the security databases in secure computer 80.

In the operational state the network connections are enabled and the Server will execute only software which has been compiled and installed as executable by an assured party.

The two states are reflected in two separate kernels. The administrative kernel is not subject to Type Enforcement. Instead, it is network isolated and accessible only to authorized personnel. This means that in administrative kernel mode, secure computer 80 cannot be seeded with malicious software by any but the people charged with system administration.

On the other hand, the operational kernel is subject to Type Enforcement. This means, for instance, that executable files stored in the memory of secure computer 80 cannot be executed without explicit execution privileges. In one such embodiment, executable files cannot be give execution privileges from within the operational kernel. Instead, secure

computer 80 must enter administrative kernel to grant execution privileges. This prevents execution of malicious software posted to secure computer 80 memory. Instead, only executables approved by operational administrators while in administrative kernel mode ever become executable within 5 operational kernel mode of secure computer 80. In such an embodiment, administrative kernel can be entered only from either a manual interrupt of the boot process to boot the administrative kernel or by booting secure computer 80 from a floppy that has a pointer to the administrative kernel.

These restrictions provide the following advantages:

- 10 - Defense in Depth: If an attacker should find a vulnerability in a system 40 subsystem, the damage that attacker can cause is limited to that subsystem. This prevents well-known attacks where a vulnerability in, e.g., the mail subsystem can be exploited to take over an entire installation.
- 15 - Silent Alarms: The Type Enforcement supersedes and constrains the traditional "root" and "superuser" privileges of insecure Unix. Attempts to exercise these privileges in system 40, or to violate other constraints of Type Enforcement, result in alarms being raised in administrative processes. No signal or indication of attack detection need be given, however. Instead, system 40 can, if desired, gather 20 data to trace the source of the attack, feed false or misleading data to the attackers or take other appropriate countermeasures.
- Open Security Architecture: The modular design means new Internet services can be provided quickly and securely.

25 An example of an assured pipeline appears in the diagram shown in Fig. 5a. The flow of data between processes in Fig. 5a is controlled by the access enforcement mechanism of the Intel Pentium processor. Virtual memory translation circuitry within the Pentium processor includes a mechanism for assigning access privileges to pages of virtual memory. This 30 ensures that control is imposed on every fetch from, or store to, the machine memory. In this way, the protection is made continuous. The Pentium access control mechanism enforces the following modes of access:

- Read Only (R): Data values may be fetched from memory and used as inputs to operations, but may not be modified or used as program text.
- Read Execute (RE): Data values may be fetched from memory and used as inputs to operations, and may also be used as program text, but may not be modified.
- Read Write (RW): Data values can be fetched from memory and used as inputs to operations, and may also be stored back in modified form.
- 10 • No Access: The data cannot be fetched from memory for any purpose, and it may not be modified.

The diagram in Fig. 5a then shows how these hardware-enforced accesses are used to force data flowing from internal network 82 to the Internet to go through a filter process, without any possibility that the filter is bypassed or that filtered data is tampered with by possibly vulnerable software on the Internet side of the filter.

The access a process has to a data object via Type Enforcement is defined by an entry in a central, protected data structure called the *Domain Definition Table* (DDT). A representative DDT is shown in Fig. 5b. A 20 Domain name denotes an equivalence class of processes. Every process in execution has associated with it two Domain names which are used to control its interaction with object and with other Domains. The real Domain of a process is used to control Domain to Domain interactions and to grant or deny 25 special, object-independent privileges. The effective Domain of a process is used to control its access to objects. The real and effective Domains of a process will generally be identical; the circumstances in which they differ are described below.

A Type name denotes an equivalence class of objects. Objects 30 are, in general, the "base types" of BSD/386 Unix: files, directories, etc. There are eight default subtypes: *file*, *directory*, *socket*, *fifo*, *device*, *port*, *executable*, and *gate*. The implied default subtype *pipe* is, in effect, untyped

- The effective Domain of the process requesting the access or action.
- The creator field of the object Type.
- The subtype field of the object Type.

5 The result of "indexing" is the retrieval of a set of access attributes. The term "attribute" is used instead of "mode" because some of the attributes define immediate side effects. The selection of attributes was governed by the following considerations.

- To constrain the modes of access which processes may exercise on objects.
- To prevent the execution of any application software other than that which has been installed through the controlled administrative environment.
- To enable the spoofing of attackers so that the attack response facilities can be used to trace them at the physical packet level. This required a more sophisticated response to illegal accesses than just shutting down the offending process.

10 The possible access attributes and their meanings are given in the table in Fig. 7.

15

20

Interactions Between Domains and Domains

The rules which govern the setting of the real and effective Domains of a process are as follows:

- Processes which are created by a *fork* syscall have their real and effective Domains set to the real and effective Domains of the parent process.
- If the executable used by *execve* syscall is of subtype *exec*, the real and effective Domains of the process are unchanged.
- The *makedomain* syscall may be used to change the real Domain of a process at the same time the executable is changed (analogous to *execve*). The new real Domain must be allowed

25

30

by the DIT (the process is as shown in Fig. 11), and the effective Domain is changed to the new real Domain.

- the *changedomain* syscall may be used to change the real Domain of a process without changing the executable.

5 - if the executable used by *execve* is of subtype *gate*, the effective Domain of the process is set to the creator field of the full Type name of the executable. This action is called implicit gating. The new effective Domain must be allowed by the DIT.

10 - The *gate* syscall may be used to change the effective Domain of a process without changing the executable. The new effective Domain must be allowed by the DIT. This action is called explicit gating.

- The *ungate* syscall may be used to change the effective Domain of a process back to its real Domain. This action is called 15 ungating.

Consider the case where a process running in the Mail Domain has execute access to files of Type Mail:exec and SMTP:gate. Further assume that there 20 exists a Domain MIME. Then the new and effective Domains resulting from the relevant syscalls are shown in the table in Fig. 8. Gating facilities are not absolutely necessary for Type Enforcement to work. They exist for the following reasons:

25

- To simplify the DDT, by reducing the number of Types that would have to exist simply to implement inter- Domain data flow.
- To improve performance, by reducing the amount of copying and signalling required to coordinate activities in different Domains.
- To facilitate the porting of existing code whose process 30 structure was not determined or influenced by

considerations of least privilege or confinement of effect.

Gating permits a process to temporarily become a member of another Domain. The "home" or permanent Domain of the process is called 5 its real Domain and the temporary or assumed Domain is called the effective Domain.

Implicit gating is used when it is necessary to strictly control the manner in which the effective Domain's accesses are used. Implicit gating "ties" the temporary Domain change to a specific executable which has been 10 subjected to extra scrutiny to insure that the effective Domain's accesses are used safely. The "tying" of the Domain change is done because the Domain change is a side effect of *execve*'ing a special executable: one whose subtype is *gate*. Implicit gating also allows Domain changes to be defined by changing the Type of an executable instead of inserting explicit calls into the 15 source code.

Explicit gating is used when a looser control on the temporary Domain transition is appropriate, or when the "tying" of the gating to a specific executable would require excessive restructuring of existing software.

Domain changes are controlled by the DIT. The logical 20 structure of the DIT is a table with an entry for each Domain. The logical structure of each entry is that of two pointers, one to a list of allowed real Domains and the other to a list of allowed effective Domains. Thus, if a process executed a *makedomain* or *changedomain*, the real Domain of the process selects the entry and the Domain given by the *domainname* argument 25 must be on the list of allowed real Domains for the Domain change to happen. Likewise, if a process executes a *gate*, the Domain given in the *domainname* argument must be on the list of allowed effective Domains. Finally, if a process executes an *execve* of an executable whose subtype is *gate*, the *creator* Domain of that executable must appear on the list of allowed 30 effective Domains.

Certain kernel syscalls are restricted to processes executing out of privileged Domains. In the preferred embodiment of Type Enforcement

two levels of checks are made. First, the normal BSD UNIX permissions are checked; if these permissions cause the operation to fail, the system call returns the normal error code. If the UNIX permissions are adequate, the TE privileges are checked next, (and thus in addition to the UNIX permissions).

5        The following BSD system calls have been modified to properly implement Type Enforcement. The modified calls have been grouped into four groups for ease of explanation.

10        The first group of system calls that require modification are those that set or affect the identity and/or state of the computer. Two of these system calls affect the computer's internal time: *settimeofday* and *adjtime*. Both of these system calls have been modified to require the <can\_set\_clock> privilege before the request will be honored. In the event of a privilege violation, the system call will raise an Alarm, will not honor the request, but will return success.

15        Other system calls which affect the computer's notion of self identity are *sethostname* and *sethostid*. Both of these system calls have been modified to require the <is-startup> privilege before the request will be honored. In the event of a privilege violation, the system call will raise an Alarm, will not honor the request, and will return the EPERM error flag. The 20 last system call affects the computer's runtime status, *reboot*. The *reboot* system call has been modified to require the <admin-reboot> privilege before the request will be honored. If the request is honored, the computer will boot to the admin kernel (single-user mode only with networking disabled). In the event of a privilege violation, the system call will raise an Alarm, will not honor the request, and will return the EPERM error flag.

25        The second group of system calls that require modification are those that allow interaction with the computer's filesystem. The *open* system call has been modified to become the primary TE check. After performing the normal BSD UNIX permission checks, the TE check is performed. An 30 Alarm is raised if the TE check returns null (no permissions), or if the caller asks for read but the <ddt\_read> privilege is not set, or if the caller asks for write but the <ddt\_write> privilege is not set. The *creat* system call has been

modified to set the new file's Type to <creator:file>. Additionally, the creation of a new file implies a write operation on the directory, which in turn implies that the TE-modified *open* system call will be used to open the directory file, which in turn implies that TE can be used to control the success or failure of the *creat* system call. The *unlink* and *rename* system calls are modified in like manner. The *unlink* system call requires the <ddt\_destroy> privilege. The *rename* system call requires the <ddt\_rename> privilege on the "from" file, and if the "to" file exists, it further requires the <ddt\_destroy> privilege on the "to" file. In the event of a privilege violation, both the *unlink* and *rename* system calls will raise an Alarm, will not honor the request, but will return success. The *access* system call is modified to require the <mode> privilege on the file pointed to by the path. In the event of a privilege violation, the *access* system call will raise an Alarm, will not honor the request, but will return success. The *chflags*, *fchflags* and *quotactl* system calls are modified in alike manners. All are modified to perform no functions. Attempts to call them will raise an Alarm, will not honor the request, and will return EPERM. The *mknod* system call is modified to perform no function. Attempts to call it will raise an Alarm, will not honor the request, and will return EPERM.

The third group of system calls that require modification are those concerning process creation, maintenance and tracing. The *fork* system call has been modified so that the child process inherits both the real and effective Domains of the parent process. The *execve* system call is modified to require the <ddt\_exec> privilege on the file pointed to by the path before the request will be honored. The real and effective Domain of the process remain unchanged. In the event of a privilege violation, the system call will raise an Alarm, will not honor the request, but will return success. The *ktrace*, *ptrace* and *profil* system calls are modified in alike manners. All are modified to perform no function. Attempts to call them will raise an Alarm, will not honor the request. The *ktrace* and *ptrace* system calls will return EPERM, whereas the *profil* system call will return EFAULT.

The *mprotect* system call is modified to perform no function. Attempts to call it will raise an Alarm, will not honor the request, and will return EPERM.

The fourth group of system calls that require modification are 5 those that relate processes to user ids. The *setuid* and *seteuid* and *old.setreuid* system calls are modified in alike manners. All are modified to require the <suppress\_su\_alarm> privilege before the request will be honored. In the event of a privilege violation, the system call will raise an Alarm, will not honor the request, and will return success. The *acct* system call is modified 10 to perform no function. Attempts to call it will raise an Alarm, will not honor the request, and will return EPERM. The *setlogin* system call is modified to require the <can\_setlogin> privilege. In the event of a privilege violation, the access system call will raise an Alarm, will not honor the request, but will return success.

15 A final set of system calls consists of those that are removed entirely from the BSD UNIX kernel. This set of system calls includes: *obs\_vtrace*, *nfssvc*, *asynch\_daemon*, *getfh*, *shmsys*, *sfork*, *getdescriptor*, and *setdescriptor*. (The set of system calls that were added to the BSD UNIX kernel is discussed elsewhere.)

20 The manner of searching the DDT is given in the diagram in Fig. 10. The algorithm is as follows:

- Obtain Type name 100 from the inode, where it is stored as a *long*, and parse it into two parts: the *creator* Domain  $D_C$  and the subtype name  $T_S$ .
- Obtain effective Domain 102,  $D_E$ , from the process data base. If the executable object attempting the access is of Type  $D_G:gate$ , change  $D_E$  to  $D_G$ . (Note that a previous search of the DDT must have returned *ddi\_exec* on the *exec* or *gate* object for this process to have begun.)

- If  $D_E = D_C$ , and  $T_S$  is one of the default subtypes such as *file* (but not one of the "static" subtypes *gate* or *exec*) then return *ddi\_read* + *ddi\_write* + *ddi\_rename* access.
- If  $D_E \neq D_C$ , or if  $D_E=D_C$  and  $T_S$  is not one of the default subtypes, then search the DDT structure 104 for the entry corresponding to  $D_C$ . If no such entry exists, search the structure for a "wildcard" entry. If neither an entry corresponding to  $D_C$ , or a "wildcard entry" exists in the structure, assign null access.
- If an entry for  $D_C$  exists, search the *subtype* list 106 it points to for an entry corresponding to  $T_S$ . If no such entry exists, search the *subtype* list it points to for a "wildcard subtype." If neither such entry exists, assign null access. If an entry for  $D_C$  does not exist, but a "wildcard" entry does, search the *subtype* list the "wildcard entry" points to for an entry corresponding to  $T_S$ . If no such entry exists, search the *subtype* list the "wildcard entry" points to for a "wildcard subtype." If neither an entry corresponding to  $T_S$ , nor a "wildcard subtype" exists in the *subtype* list, assign null access.
- If a *subtype* list entry for  $T_S$  exists, search the Domain vector 108 it points to for an entry 110 corresponding to  $D_E$ . If no such entry exists, search the Domain vector for a "wildcard Domain." If neither an entry corresponding to  $D_E$ , nor a "wildcard Domain" exists in the Domain vector, assign null access.
- If a Domain vector entry for  $D_E$  exists, return the access values it contains. If a "wildcard Domain" entry exists in the Domain vector, return the access values it contains. If neither a Domain vector entry for  $D_E$ , nor a "wildcard Domain" exists in the Domain vector, return null access.

The above algorithm describes the "logical" process of searching the DDT; the actual implementation is described next.

As noted above, in one embodiment, Domains and *subtypes* are stored as four printable character constants (white space doesn't count as printable - also, '\*' is excluded). Due to constraints imposed by the fact that BSDI Release 1.1 does not contain complete source code, only the first character of a Domain and the first three characters of a *subtype* are significant, and thus must be unique. Furthermore, there is a convention that *subtype* names that appear globally (i.e., both default *subtypes* and *subtypes* used by more than one Domain) be made of lowercase characters, while 10 private *subtypes* be made of uppercase characters.

These four character names are represented by C constants. For Domains, these constants begin with a D, while for *subtypes*, these constants begin with a T. The following character should also be in 15 uppercase (e.g., *DRoot*, *TFile*). There is also two special constants: *kWildcard* = \*\*\*\*, which matches any *subtype* or Domain, and *kEOL* = 0, which is used to mark the end of a list. These constants are all defined inside a list of *enum*'s since using *#define* would result in too many compiler 20 warnings (the C compiler warns about multi-character constants, by using *enum*'s, it will only warn once for a given constant).

There are six default *subtypes*, based on existing Unix types:

```
/* Here are the default types... */
TFile = 'file'
TDirectory = 'diry',
25      TSocket = 'sock',
Tfifo = 'fifo',
TDevice = 'devi',
TPort = 'port',
TExec = 'exec',
30      TGate = 'gate'
};
```

*TExec* is a special *subtype*, which can only be assigned by the isolated administrative kernel. It represents executables which any Domain can execute if execute access is allowed by the DDT. *TGate* is a special sort 35 of *TExec* - what it does is change the effective Domain in which a process is

executing to the *creator* Domain of the *gate*. It only does this if the starting Domain has execute access to the file of *subtype gate*. After "gating," a process now acts like it is in the *creator* Domain for the purposes of the DDT 5 checks only - any checks against the DIT are made with the real Domain, rather than the effective Domain. Needless to say, a *gate* is a powerful and potentially dangerous thing - just like the *setuid* bits which gating is designed to replace. Note that there is a special check in the normal DIT checks - if we are attempting to change to the real Domain, we don't bother to check the DIT of the effective Domain (like we otherwise normally would). This 10 maneuver is "ungating" - explicitly leaving the gated Domain and returning to the original Domain.

There is only one pre-defined Domain:

```
15      enum {
              DRoot = '$$SYS', /* Root is actually a special alias for
                                the zero domain that the system is started
              in */
          };
```

which is used to represent system level defaults - whenever a Domain that hasn't been explicitly set (for either a file or a process), *DRoot* is 20 used for the Domain value in permission checks.

The DDT is made of a three level table, indexed by the file's *creator* Domain, file's *subtype* and then finally by the executing Domain. This yields a set of access permissions:

```
25      typedef unsigned long ddt_permissions;
          enum {
              ddt_read = 1,
              ddt_write = 2,
              ddt_rename = 4,
              ddt_exec = 8,
              ddt_trigger = 0x10,
              ddt_chcreator = 0x20,
              ddt_destroy = 0x40,
          };
```

These permissions work mostly as expected - *ddt\_read*, 35 *ddt\_write* are for read and write; *ddt\_rename* permits changing the name of the file; *ddt\_exec* is used to grant execute permission; *ddt\_destroy* is required to delete a file. *ddt\_chcreator* is much like create permission, but since files

are created with a default *subtype*, this permission allows the given Domain to change the subtype and *creator* of the file to the corresponding *subtype/creator* pair. *ddt\_trigger* isn't really a permission - rather, any checks to this specific file will automatically trigger an alarm, regardless of what permission is asked 5 for or granted. This allows, for example, a "reverse trojan" file that would never be executed except by an attacker, in which case an alarm would be triggered and packet-level auditing performed.

The indexing begins with an array "indexed" by Domain:

```
10  typedef struct {
11      type_name src_domain; /* What domain this entry is for */
12      unsigned long domain_flags; /* The "global" permission flags
13      */
14      type_name * the_ddt; /* What domains we can enter into */
15      ddt_type_list the_ddt; /* The permissions for our types */
16  } permission_table;
```

This array should have an entry for every Domain. For the DDT, this table is searched until the *src\_domain* matches the *creator* Domain 20 of the file. Assuming that it is found, we then look at the *ddt* an array "indexed" by *subtype*:

```
/* This is the permission for a specific domain, listing all its types */
25  typedef struct {
26      type_name the_type; /* The subtype */
27      ddt_domain_vector_entry * the_vector; /* A list of what can be done to
28      it */
29      ddt_type_list_entry, *ddt_type_list;
```

30 We just look through this list until we either find the *subtype*, a wildcard, or the end of the list (in which case we return no permission). We then need to look at the appropriate *the\_vector* - an array "indexed" by Domain:

```
35  typedef struct {
36      type_name the_domain; /* The using domain */
37      ddt_permissions the_permission; /* What it can do */
38  } ddt_domain_vector_entry, *ddt_domain_vector;
```

This is search for the executing Domain, and if found, we return the permission, which contains the flags for that access.

Searching the DIT starts like searching the DDT. We look through the global table for the starting Domain, and find the appropriate list of Domains the\_ddt. This is simply a list of Domains, terminated with kEOL.  
5 We search through that list, and if we find the desired destination Domain, we can make the transition to it.

Every Domain also has a list of privileges that it can perform:

```
10 enum {
11     can_ch_type = 0x00010000, /* We can call ch_type, changing
12     type */
13     suppress_su_alarm = 0x00020000, /* Allow process to think it is
14     su */
15     admin_reboot = 0x00040000, /* Allow reboot */
16     can_set_clock = 0x00080000 /* Can set the clock */
17     can_setlogin = 0x00100000 /* Can perform setlogin */
18     is_startup = 0x00200000 /* can perform startup actions */
```

20 We look through the permission table to find the appropriate Domain, and then get these permissions from the appropriate domain\_flags field. Note that there is no explicit "can\_ch\_domain" permission; restrictions on Domain transitions are enforced by the DIT.

25 Since each and every array must be a separate C structure, every array needs to have a unique and meaningful name to connect one array to its parent. This is best explained in a "simple" example.

```
/* NB: In the initial implementation, domains need to have unique
   first characters */
30 enum {
31     DRoot = '$SYS', /* Root is actually a special alias for the zero
32     domain that the system is started in */
33     DUpdate = 'sync',
34     DSwap = 'Swap',
35     DUserSession = 'User',
36     DSyslogd = 'Logd',
37     DCron = 'Cron',
38     DRouted = 'Rout',
39     DSendmail = 'mail',
40     DIntd = 'inet',
```

```

5      DTelnet = 'inet',
      DShell = 'rshd',
      DRexec = 'exec',
      DFinger = 'fing',
      DNetwork = 'xnet',
      DLpd = 'lpd.',
      DPortmap = 'port',
      DFsck = "Fsck",
      DQuota = 'quot',
      DXDOS = 'dosx',
      DFTP = 'Tran',
      DInnd = 'News'
10
15  };

```

15 These are just a list of sample Domains:

```

type_name Root_dit[] = {
20      DUserSession, DSyslogd, DUpdate, DCron, DRouted, DLpd,
      DPortmap,
      DSendmail, DInetd, DInnd, kEOL

```

These are some addition private *subtypes* for our example:

```

25      /* Here are some other types */
      enum {
30          TStartup = 'Stup',
          TConfig = 'Conf',
          TCronJobs = 'CJob'

```

30 This is the list of Domains that the root Domain can change to.  
 The naming convention here is *DomainName\_dit*, where *DomainName* is the name of the constant for that Domain without the leading "D". The Domain list is terminated with a *kEOL*.

```

35      { ddi_domain_vector_entry Root_Startup[] = {
      { kWildcard, ddi_read },
      { kEOL } };

```

40 This is our first Domain vector. The naming convention is *CreatorDomainName\_TypeName*, where *CreatorDomainName* is the name of the constant for the creating Domain (without the leading "D"), and

*TypeName* is the name of the constant of the *subtype*. Vector is initialized to contain a list of Domains and permission pairs, terminated with { *kEOL* }.

```
5      ddi_domain_vector_entry Root_default [] = {
        /* This is the default permissions for all procs on all unassigned files
       */
       { kWildcard, ddi_read|ddi_write|ddi_rename },
       { kEOL }
    };
10
```

*Root\_default* will be the Domain vector for creator *Root*, and subtype *KWildcard* - basically the default for any subtypes created by *DRoot* who otherwise wouldn't have a special Domain vector.

```
15     ddi_domain_vector_entry Root_Exec[] = { /* Default for executables
      */
       { kWildcard, ddi_exec|ddi_read },
       { kEOL }
    };
20
```

Another Domain vector, this time for all executables owned by the system.

```
25      ddi_type_list_entry Root_types [] = {
        { TStartup, Root_Startup },
        { TConfig, Root_Startup },
        { TExec, Root_Exec },
        { kWildcard, Root_dDefault },
        { kEOL }
    };
30
```

Once we have all the Domain vectors for a given creating Domain, we can make the corresponding *subtype* list. The naming convention is *CreatingDomain\_types*. It is composed of pairs of *subtypes* and the corresponding (previously declared) Domain vectors. Note that it is possible for more than one *subtype* to use the same Domain vector (in this case, both *TStartup* and *TConfig*).

```
35      permission_table Rover [] = {
        { DRoot,
          can_ch_type | can_ch_creator | admin_reboot |
          can_set_clock,
          Root_dit, Root_Types },
40
```

```
    {DInetd,
     0,
     Inetd_dit, NULL }
     { kEOL }
5    };
```

Here is the master permission table "Rover". It is composed of a list of Domains (two in this case). Each entry contains the Domain name, the permissions for that Domain, its DIT and its *subtype* list. If the DIT is NULL, then no Domain transitions out of that Domain are allowed. If its 10 *subtype* list is NIL, then there is null access to all *subtypes* of that creating Domain. The last entry, of course is the *kEOL* termination.

Every process runs in a Domain, which is stored in the kernel *proc* data structure. This property is copied to processes that are forked, and is unchanged by executing most binaries and shell scripts. The Domain can 15 be explicitly changed via the *makedomain* system call, which, if permitted, changes Domain for that process from that point forward. Privileges of a given Domain can also be granted to something running in another Domain via a "gating" process - a process that executes a file of *subtype gate* will, assuming there is execute permission granted to the current Domain for that 20 file, temporarily assumes the privileges of the *creator* of the *gate* file. This is accomplished by an "effective Domain:" field in the kernel *proc* data structure. This field is also copied during forking, and is reset when *makedomain* is successfully called (reset to the new Domain specified). Most importantly, the effective Domain field is used to check file access 25 permissions, but real Domain is used for checks from *makedomain*. There is, however, a special addition to *makedomain* for the purpose of "ungating" - if the process is calling *makedomain* with the real Domain, it automatically succeeds (thus resetting the effective Domain to the real Domain), allowing a process to return to the Domain that it started in. The Domain transition 30 permissions are all handled in *domain\_to\_domain*. This routine first looks up the source Domain in the permissions table. It will use the *kWildcard* entry, if any, to provide default source Domain permissions. It then looks in the DIT vector for the destination Domain, and, if found, allows the transition. It

will not, however, use a wildcard in that vector, since this would allow a given Domain to transition to every other Domain.

The most important check that *execve* makes is to check for *ddi\_exec* access. It looks at the *subtype* and *creator* of that which is to be executed and the effective Domain of the current process (not the real Domain), and makes sure that there is *ddi\_exec* access. If there is it also compares the *subtype* of the file to see if it is *gate* - if so, we change the effective Domain to the *creator* of that file.

There is also logic in *execve* that makes sure that we don't gate 10 by mistake - the old effective Domain is grabbed at the start of *execve*, and any time that an error is returned, we first restore the old effective Domain.

*chtype/fchtype* are used to change the *subtype* and/or *creator* of a file. Because of this power, they must be carefully controlled. One of the first constraints on *chtype* is that it can either change both the *subtype* or 15 *creator*. We can never change anything about a file that we aren't currently the creator of. Furthermore, since *exec* and *gate* are special static *subtypes*, we can never make or unmake an *exec* or *gate*. This is only done from the administrative kernel. The final special rule is that we can only change to a *subtype/creator* that already exists (this is to prevent making "orphaned" 20 object, but with the special *kWildcard* type we could still specify access permissions for these things, so this rule could be removed). Note that this "check for existence" will accept wildcards in the permission table as matching whatever we pass in.

The other checks made by *chtype/fchtype* are checks to the 25 permission table. First off, the executing Domain needs to have *can\_ch\_type* permission. Then, if we are only changing the subtype of an object that we created (and all the checks in the previous paragraph pass), then we just go ahead and do that. If however, we are changing the *creator* as well, we check the *ddi* to see if our effective Domain (as opposed to real Domain - see gates 30 for more detail) has *chcreator* capabilities for the *creator/subtype* that we are going to change the file to (we already know that we created it, so we don't care what the *subtype* is). If we do then we change it, if not, then we don't.

Actually changing the *subtype*, since we are hacking *subtype* and *creator* into the flags field of the *vnode*, requires us to be running as root (since we are changing both words of the flags field, and *VOP\_SETATTR* seems to care). So, before calling *VOP\_SETATTR*, we first save the *cr\_uid*, 5 set it to zero, and then restore it. When we modify *VOP\_SETATTR* to write our *subtype* and *creator* to the real places in the inode, this will be removed.

*check\_ddt* takes an effective Domain and a *creator:subtype* pair and looks for specific access attributes, returning those that correspond to permissions, and raising alarms if things don't work as expected. The first 10 thing that *check\_ddt* does, after mapping any potentially undefined fields to *DRoot* and/or *TFile* (if the *subtype* or *creator* is zero, such as on a file system not properly set up), is check for the default subtypes. If the source Domain is the same as the *creator*, and the subtype is one of the default eight 15 *subtypes*, the returned access attributes are *ddt\_read* + *ddt\_write* + *ddt\_rename*.

Otherwise, we need to look up the *creator:subtype* in our tables. If we find them (or appropriate wildcard matches), we then search the Domain vector to find the source Domain. If we find that (or again, the wildcard), the 20 return permission is taken from there. If we never find one of the respective entries, the return permission is no permission.

The last step in *check\_ddt* is to see if the return attribute is inconsistent with the permissions asked for by the caller, or if the resulting permission includes the *ddt\_trigger* attribute. If either of these cases are true, then we need to log this request to the alarm mechanism. This involves 25 writing out the process id, the name of the file, the parameters and what permission is returned. The alarm processing would, at that point, take appropriate action.

In addition, the system 40 shown in Fig. 4 is constructed so that no software may be loaded into it except under the control of the System 30 Administrator, and even then only when the system is disconnected from all networks. (This is a function of the two kernels: operational and administrative, as described above.)

The Type Enforcement mechanism allows a strict *least privilege* design to be defined and enforced. Least privilege is a way of achieving confinement, or the limiting of a software module's effects. A least privilege design is one in which software only touches the data it needs to operate and

5 only touches it in ways that the designer intended. Unwanted side effects, whether from bugs or malicious trojan horses, are then limited to the module's "immediate vicinity." This fundamental ability of Type Enforcement, when properly applied, stops dead the most common types of attacks, where a vulnerability in one application is used to interfere with, or take control of,

10 more critical sections of the system.

In order to take advantage of this capability, the application only needs to follow traditional Unix practices and be implemented as several processes. These processes can be assigned to a distinct class, as can the data that they access. The DDT can be configured to allow only the least amount

15 of access necessary for the desired functionality.

The Type Enforcement described above permits a security architect to construct a set of interconnected applications and protect them with countermeasures such as data filters. The architect can do this with the confidence that the applications and countermeasures will be isolated from

20 each other and share data only in the ways the architect defines. This enables the architect to upgrade system 40 quickly to respond to changes in threat, by adopting new countermeasures; to secure new applications, by constructing countermeasures that address the specific vulnerabilities of the application; and to implement customer-specific security policies which balance risk

25 against operational effectiveness.

Since Type Enforcement defines pipelines and subsystems which are independent with regard to privilege, the addition of a new subsystem or the extension of a pipeline does not, in and of itself, obsolete the assurance evidence produced for the previous structure. Rather, the

30 assurance team can examine the new interactions and decide precisely which conclusions about isolation are still valid and which have to be re-examined.

Type Enforcement has also demonstrated its ability to support cryptography, whether implemented in hardware and software. Cryptographic processing, with its requirements for separation of plaintext and ciphertext, is inherently a pipelined process. This is true whether the cryptography is 5 placed in its traditional "inline" position or whether it is used in the "coprocessor" mode required for the more advanced services such as digital signatures and non-repudiation.

Type Enforcement is better than the basic Unix protection mechanisms for two reasons: it is centralized instead of decentralized, and it 10 does not permit any process to have global, uncontrolled access. In Unix, individual programs use the setuid mechanism to set their own privilege level. A particular privilege level, called "root," or "super-user," lets a user do anything they want to the system: observe and manipulate data, disable auditing, install trojan horses, or masquerade as other users. This combination 15 of decentralization and potential global privilege is deadly. Decentralization means that there is no one place you can look to see if the system is configured securely. Global privilege means that a single vulnerability or configuration mistake can be catastrophic.

Type Enforcement eliminates both these problems. If you stop 20 a system 40 as described in Fig. 4 and dump the DDT you can tell for sure which code could ever have touched which data. You can never tell that in a Unix system. And nobody ever gets global privilege when secure computer 80 is attached to a network.

In the preferred embodiment, the Type Enforcement restrictions 25 supplement, but do not replace, the standard Unix permissions. That is, you can set Unix permissions to give less, but not more, access than Type Enforcement allows. And super-user privilege is still there, but it cannot be used to exceed the Type Enforcement limitations.

In one embodiment, a system 40 detects an attack in progress 30 (as a result, for instance, of a Type Enforcement violation) it trips a "silent alarm" which is responded to by application-specific countermeasure software.

This software can, depending on the nature of the attack, do the following things:

- Capture the IP address of the attacking site, enabling calls to site administrators to trap attackers in the act.
- 5     • Feed the attacker false and misleading data.
- Feed the attacker useless but "interesting" data so he stays online and can be traced.
- Feed the attacker data containing covert identification data that can be used to prove that data was stolen from this site.

10

In one embodiment, a binary filter is used to ensure that neither executables nor encrypted files are transferred into or out of system 40. (The prohibition against executables is an attempt to capture malicious software transferred into the system and to detect the posting of potentially proprietary object code from system 40 onto the Internet. The prohibition against transfer of encrypted files is an attempt to prevent the posting of encrypted versions of proprietary information either to or from system 40.) In one binary filter embodiment, text is analyzed to determine if it is written in English. The filter looks each character and its next neighbor and determines the frequencies of pairs of letters ("a diagraphic index of correlation"). If the index of correlation approximates what would be expected for English text, the file is probably English text and can be transferred. If not, filter 92 stops the transfer.

25 Operation of the Secure Wide-Area Access System

When a Client desires to put information out on Public Network 74, he or she must first use the Local Cryptography to establish and authenticated and protected interaction with Secure Computer 48. The Client then issues the requisite commands through the Client interface, and these 30 commands and their associated are then executed and controlled by the integrated set of services and filter counter-measures on the Secure Computer. The Public Network Protocol and Cryptography module then selects the

appropriate authentication and protection mechanism for the interaction on Public Network 74. Depending on the protocols and cryptography used, Public Network 74 and Cryptography module 70 may then perform cryptographic and format transformations on the data. Most commonly, these would involve 5 decrypting data that was encrypted using Local Cryptography, changing its format from a local messaging or data transfer format to a global standard, and encrypting using Global Cryptography. At the same time, Secure Computer 48 can generate an audit record and protect it with cryptographic keying material accessible only to authorized administrators.

10 If authentication is required, Secure Computer 48 can either "endorse" or "notarize" the data using cryptographic keying material of its own, or it can act as a secure storage and selection facility whereby the local authentication of the Client is used to select the personal keying material used to authenticate the Client's identity on Public Network 74. Secure Computer 15 48's facilities can use other information, such as time of day, content of the data, etc., as well as the facilities of the Local Cryptography to decide whether or not to perform authentication of the outbound information.

20 An important special case is where two systems 40 at two different sites belong to the same organization. Such a situation is shown in Figs. 12 and 13. In Fig. 12, two systems 40 are connected by an external Public Network 74. In Fig. 13, two systems 40 connected by an external Public Network 74 can also communicate with an unclassified workgroup 100 or with individual computers 102 and 104 connected directly to Network 74. In such cases, special protocols and keying material can be used to identify 25 the systems to each other and indicate special actions, such as administrative changes and alarms. In addition, systems 40 can easily distribute keys between themselves in a secure manner. In one embodiment, systems 40 include Trusted Path software which can be used to establish a trusted path between independent systems 40 over Public Network 74.

30 Inbound information flow is essentially symmetric to outbound: the data is received from Public Network 74, if necessary decrypted and has its authentication checked, and then is passed through the Filter

Countermeasures 68 to determine whether the organizational security policy allows data of that label, format, or content to be released into Private Network 64. If it does, Secure Computer 48 uses Local Cryptography to protect and authenticate the transmission to Client Workstation 63. When the 5 Client accesses the data, he or she can use that cryptography to verify that the data is what it was authenticated to be over the Public Network 74.

#### Advantages Over Other Methods of Securing Data Transfer

The general advantages of the invention derive from its 10 centralization of security services in Secure Computer 48. This centralization takes advantage of the fact that Client workstations 63 must be supported by centralized services such as directories for electronic mail, databases of security attributes, and archival storage of cryptographic keys. Thus every Security Architecture which makes use of cryptography is, to one degree or 15 another, centralized.

Similarly, the facilities for detecting and responding to security alarms are most usefully centralized. Notifying a Client in a possibly exposed location that a network is possibly under attack can be counterproductive: the Client may not be authorized for such information, and even if authorized the 20 individual may not have a secure means of communicating this information to administrators. Also, one does not want to notify a possible insider threat that an attack has been detected. Thus again a degree of centralization in the architecture is unavoidable. Further centralization of security mechanisms adds both security and economic benefits:

25 1) Mechanisms at the workstations can be implemented as software and minimal, if any hardware. This implementation strategy limits the strength of the workstation mechanisms, and is only acceptable when they are "backed up" by the strength and facilities of a central Secure Computer and the restricted access inherent in a Private Network.

30 2) Concentration of the security requirements and facilities in the Secure Computer enables that unit to undergo scrutiny to a degree that would not be feasible for individual workstations. If the Secure Computer is

properly engineered it should be able to support multiple generations of workstation technology, thereby spreading the cost of specialized security engineering over time.

3) Concentration of countermeasures in a specially-engineered

5 Secure Computer raises the effort and risk of technical attacks because it forces the attacker to either reverse engineer and implement, or obtain through other means an up-to-date copy of the Computer and all its associated countermeasures software. This is harder than obtaining an instance of a workstation and its associated software. Concentration also simplifies the

10 process of responding to new or unanticipated attacks, as there are fewer units to change and those units are already under the control of security administrators.

4) Concentration also simplifies the process of administering the security databases and increases the speed and reliability with which

15 privileges can be granted and, more importantly, revoked.

5) The Secure Computer will, by its very nature, have the features which make it a near-optimum platform for key management and distribution: strong authentication of individuals, secure storage of data, controls on access to that data, and strong resistance to attacks by malicious

20 software.

6) The Secure Computer, by virtue of its central role and close interaction with security administrators, provides a logical and effective location for the receipt and response to security alarms. This characteristic combines with the ability to respond to new attacks by upgrading a smaller number of central sites and the speed and effectiveness of changes to security data bases to make the centralized approach inherently more responsive than architectures without a central point of security enforcement, where alarms, changes to software, and changes to data bases must propagate over a larger number of user-administered workstations.

30 In particular, the invention provides superior client authentication over methods such as Workstation Cryptography. In Workstation Cryptography, Clients authenticate themselves at vulnerable

workstations by means of personal identifiers such as passwords, passphrases, Personal Identification Numbers, or token-based authenticators. There is no protected backup or contextual check possible on such authentication actions; once authenticated, the Client is granted, in effect, full access to the Public Network. By contrast, the Secure Computer can keep a protected record of Client actions, and assess the propriety of an authenticated action based on that data as well as other criteria such as time of day, whether it is a business day or a holiday, or other checks of arbitrary sophistication. Conversely, the invention permits the sending of "official" data or transactions in which the 10 identity of the initiating individual is shielded from the Public Network and only the organizational identity is authenticated. This facility is useful when the nature of the transaction or data could make the Client open to unwanted attention, harassment, or retaliation.

The invention provides an advantage over Workstation Cryptography in that it is possible to enforce sophisticated, content-based organizational security policies. Such enforcement is not possible when data is enciphered at the workstation and then sent directly to the Public Network. In addition to enforcing content-based policies, the invention permits auditing of data contents to deter abuse of the privilege of sending data to the Public 20 Network. Both of these facilities are useful in countering insider threats.

The invention is superior to Workstation Cryptography in that it can handle a multitude of communications protocols and cryptographic methods without making that diversity visible at the Client workstation. This not only reduces the amount of hardware and software mechanism at the 25 multiple workstations, but it permits a single Client Interface to be used to access a heterogeneous Public Network. The Secure Computer, after it has decrypted data that was protected and authenticated by the Local Cryptography, can consult internal tables, directories on the Public Network, or the destination node to determine or negotiate a common protocol and 30 cryptographic method. All of this can be done without Client knowledge or intervention.

The invention is superior to Workstation Cryptography in that it provides a safer and more reliable framework for the management of keying material. This advantage obtains irrespective of whether secret-key or public-key cryptography is applied. The Secure Computer provides a central 5 site for the distribution and administration of keying material for all the Clients on the Private Network, and relieves the Client workstations of the responsibility of obtaining Public Network keying material for every interaction with that network. The distribution of Public Network keying material through the Secure Computer permits greater security in that the 10 identities of the requesting Clients can be hidden from the Public Network keying material service. The invention also provides superior solutions to the problems of revocation, emergency rekey, and travelling user.

The use of the Secure Computer as the central point for the distribution and administration of keying material permits the effective and 15 efficient revocation of access to either the Private or the Public Networks. In the most common configuration, secret-key methods will be used by Local Cryptography and public-key methods will be required for Global Cryptography. If the private key of a Client's public-key material are distributed to Client workstations, or, worse, stored on removable tokens that 20 the Client can remove, then revocation of the ability to decrypt (or, more importantly, authenticate) data requires a time-consuming and unreliable "broadcast" of revocation requests to all possible destinations on the Public Network. If the private key is kept on the Secure Computer, then access to it can be revoked simply and quickly.

25 The invention is superior to Workstation Cryptography in providing emergency rekey service, especially when public-key methods are used on the Public Network. If the private key part of a Client's public-key material is lost or destroyed, the Client loses the ability to decrypt data which was previously encrypted with the corresponding public key. It is not 30 sufficient to issue a new private/public pair, because there may be data in transit or in archives that was enciphered with the public key that corresponds to the lost private key. The problem then is one of saving a copy of the

private key in a highly protected fashion, and making it available only after proper authorization has been obtained. This is a natural task for a Secure Computer with protected storage and mechanisms and access limited to authorized administrators. If the organization has Secure Internetwork Services 5 Systems at multiple sites, then they can cooperate by maintaining backup copies of critical keying material for each other.

The invention is superior to Workstation Cryptography in that a Secure Computer at one site can forward the necessary keying material to another site, whether it be a Secure Internet Services System or some other 10 node on the Public Network. This forwarding can be closely controlled and audited, and the superior revocation facilities used to place a limit on the period during which the forwarded material can be used.

The invention is superior to Network Cryptography in that it permits controls, auditing, protection, and authentication to the granularity of 15 the individual Client rather than just to the node.

Although the present invention has been described with reference to the preferred embodiments, those skilled in the art will recognize that changes may be made in form and detail without departing from the spirit and scope of the invention.

What is claimed is:

1. A secure wide-area access system, comprising:
  - 5 a secure computer (48);
  - an internal network (64); and
  - a workstation (63) connected across the internal network to the secure computer;

wherein the secure computer comprises an internal network interface (66), a public network interface (72), public network program code (70) used to communicate through the public network interface to a public network, private network program code (66) used to communicate through the internal network interface to the workstation and security policy program code (68) for enforcing a Type Enforcement security mechanism to restrict access of a process to data.
- 15 2. The system according to claim 1 wherein the security policy program code comprises program code for hardening a UNIX operating system.
- 20 3. The system according to claim 1 wherein the security policy program code comprises program code for hardening a UNIX operating system, wherein the program code for hardening the UNIX operating system comprises kernel code for enforcing Type Enforcement via an operational kernel.
- 25 4. The system according to claim 1, wherein the workstation includes a client subsystem (60) for secure transfer of data from the workstation to the internal network interface of the secure computer.
- 30 5. The system according to claim 4 wherein the client subsystem is a software program running on the workstation.

6. A method of protecting a computer system connected to an unsecured external network, wherein the computer system includes a plurality of workstations connected to a private network, the method comprising the steps of:

5 providing a secure computer, wherein the secure computer comprises security policy program code for enforcing a Type Enforcement security mechanism to restrict access of a process to data;

connecting the Type Enforcement based secure computer to the private network; and

10 establishing an assured pipeline for the transfer of data and programs between the private network and the external network through the secure computer, wherein the step of establishing an assured pipeline includes the steps of:

- 15 i) placing processes within domains, wherein the step of placing processes within domains includes the step of assigning processes received from the external network to an external domain;
- ii) assigning types to files; and
- iii) restricting access by processes within the external domain to certain file types.

20

7. The method according to claim 6, wherein the step of placing processes within domains includes the steps of:

defining a domain definition table within the secure computer;  
assigning a domain name to each domain; and  
25 creating an entry for each process in the domain definition table,  
wherein the step of creating an entry includes the step of associating a domain name with each entry.

8. A secure server for use in controlling access to data stored within an internal network, comprising:  
30 administrative and operational kernels, wherein the operational kernel includes security policy program code for enforcing a Type

Enforcement security mechanism to restrict access of a process received from the external network to data stored on the internal network; and  
wherein the administrative kernel is restricted to execution only while isolated from the internal network.

5

9. A secure server, comprising:  
a processor (80);  
an internal network interface (90), connected to the processor, for communicating on an internal network; and

10 an external network interface (96), connected to the processor, for communicating on an external network;

wherein the processor includes server program code (92) for transferring data between the internal and external network interfaces and security policy program code for enforcing a Type Enforcement security mechanism to restrict access of a process received from the external network to data stored on the internal network.

10. The server according to claim 9, wherein the processor further includes encryption means for encrypting data to be transferred from the internal network to the external network.

25 11. The server according to claim 9, wherein the processor further includes filter program code for filtering data transferred between the internal and external network interfaces.

12. The server according to claim 9 wherein the processor further includes:

means for selectively filtering messages received from the internal network according to a first predefined criteria; and

30 means for selectively filtering data received from the external network according to a second predefined criteria.

13. The server according to claim 9, wherein the processor further includes formatting program code for changing the format of data transferred between the internal and external network interfaces.

5 14. A system for secure internetwork communication across an external network, the system comprising:

first and second internal networks (64);

first and second secure computers (48) connected to the external network, wherein the first and second secure computers include:

10 an internal network interface (60); and

an external network interface (72) for secure transfer of data from the first secure computer to the second secure computer over the external network, wherein the external network interface includes means (70) for encrypting data to be transferred from the first secure computer to the second secure computer;

15 a first computing system (63), wherein the first computing system includes a first client subsystem (60) connected over the first internal network to the internal network interface of the first secure computer, wherein the first client subsystem includes means for secure transfer of data between the first computing system and the first secure computer; and

20 a second computing system (63), wherein the second computing system includes a second client subsystem (60) connected over the second internal network to the internal network interface of the second secure computer, wherein the second client subsystem includes means for secure transfer of data between the second computing system and the second secure computer.

25 15. The system according to claim 14 wherein the first secure computer is a multilevel secure computer capable of recognizing data of varying sensitivity and users of varying authorizations.

16. The system according to claim 14 wherein the first secure computer is a type enforcing secure computer capable of recognizing data of varying sensitivity and of limiting access to data based on both user access rights and process access rights.

5

17. The system according to claim 16 wherein the first secure computer further comprises:

means for selectively filtering messages received from the second internal network according to a first predefined criteria; and

10 means for selectively filtering data received from the external network according to a second predefined criteria.

18. A secure computing system, comprising:

an external network (74);

15 first and second secure computers (48) connected across the external network, wherein the first and second secure computers comprise encryption means (70) for encrypting and decrypting data transferred between said first and second secure computers;

an internal network (64); and

20 a workstation (63) connected across the internal network (64) to said first secure computer (48), wherein the workstation includes means for encrypting and decrypting data transferred between said workstation and said first secure computer.

25 19. The system according to claim 18 wherein the first secure computer further comprises:

means for selectively filtering messages received from the internal network according to a first predefined criteria; and

30 means for selectively filtering data received from the external network according to a second predefined criteria.

20. The system according to claim 18 wherein the first secure computer is a multilevel secure computer capable of recognizing data of varying sensitivity and users of varying authorizations.

5 21. The system according to claim 18 wherein the first secure computer is a type enforcing secure computer capable of recognizing data of varying sensitivity and of limiting access to data based on both user access rights and process access rights.

10 22. A method of transferring data between a first and a second network connected by an external network, wherein the first network comprises a first workstation connected to a first secure computer server and wherein the second network comprises a second workstation connected to a second secure computer server, wherein each secure computer server 15 comprises a trusted subsystem, first encryption means for encrypting and decrypting data transferred between the secure computer server and its respective workstation and second encryption means for encrypting and decrypting data transferred between the secure computer server and the external network, the method comprising the steps of:

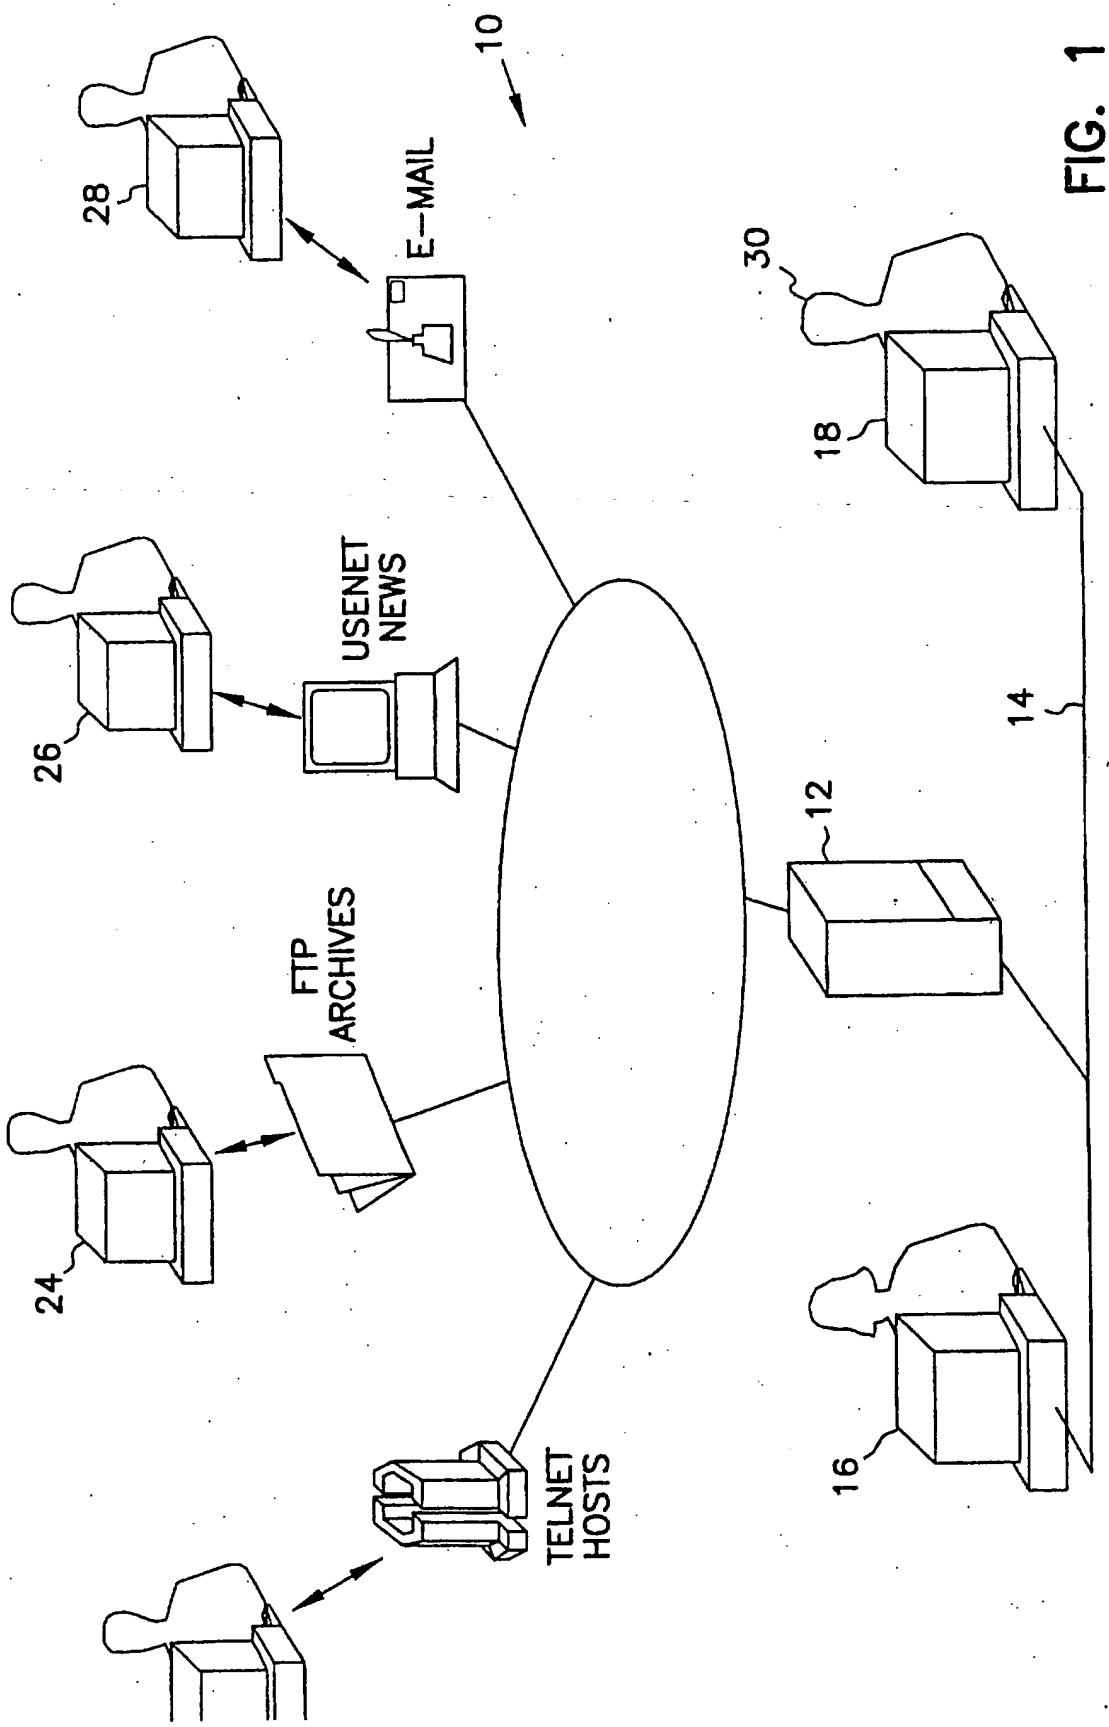
20 establishing an authenticated and protected interaction between the first workstation and the first secure computer server;

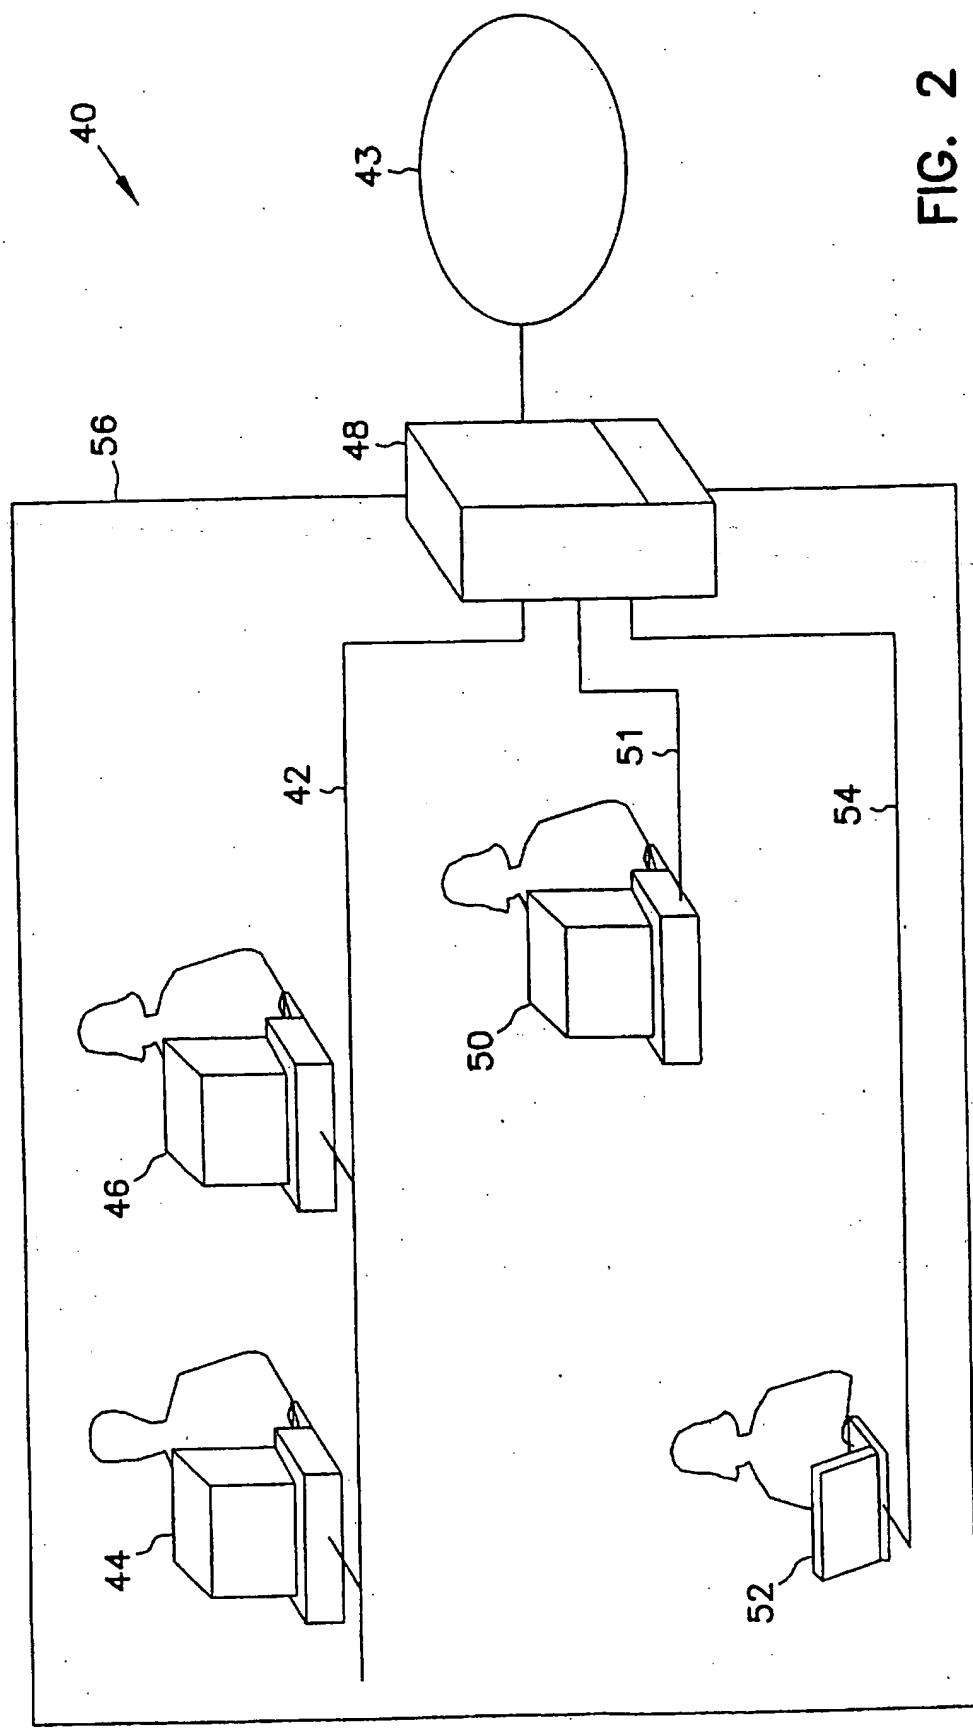
25 sending data from the first workstation to the first secure computer server;

selecting an authentication and protection mechanism for interaction on the external network;

encrypting, via the second encryption means of the first secure computer server, the data received from the first workstation; and

sending the encrypted data over the external network to the second secure computer server.





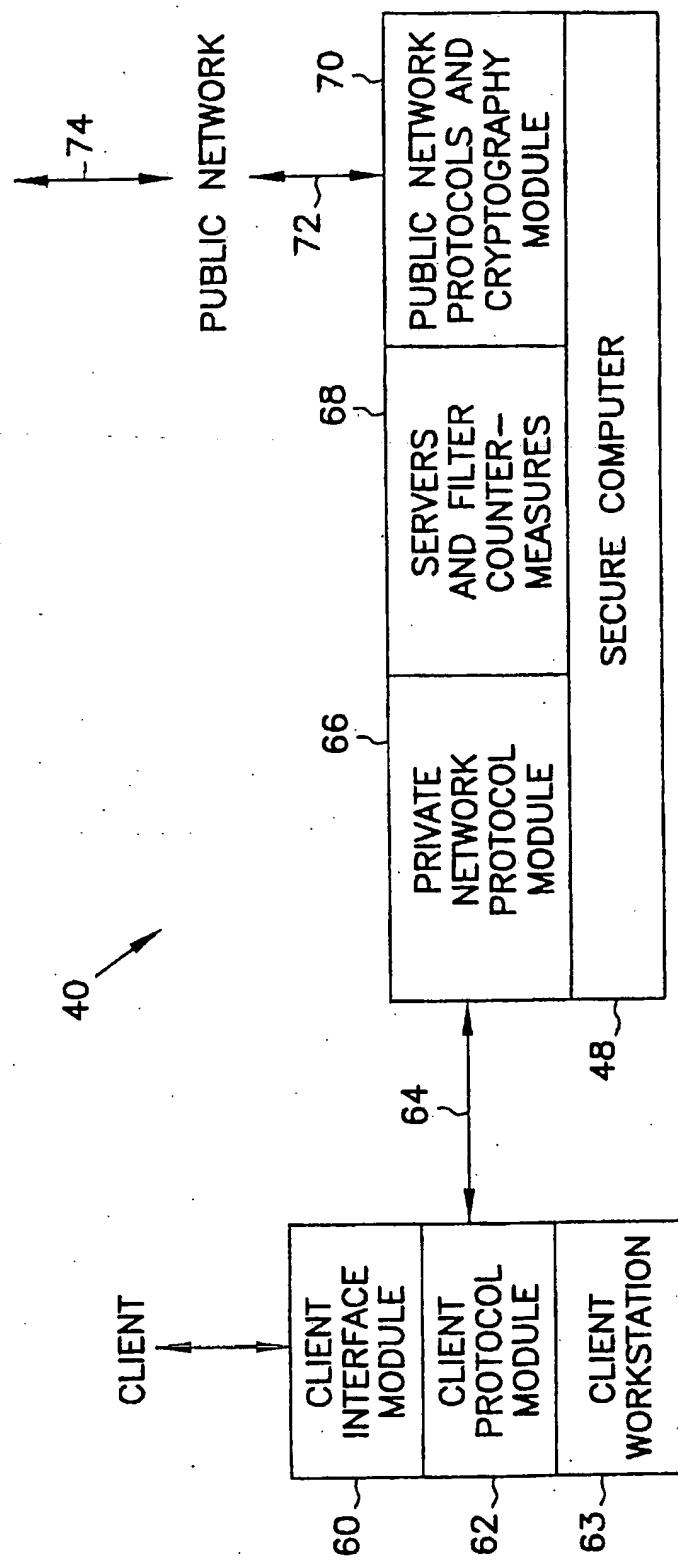


FIG. 3

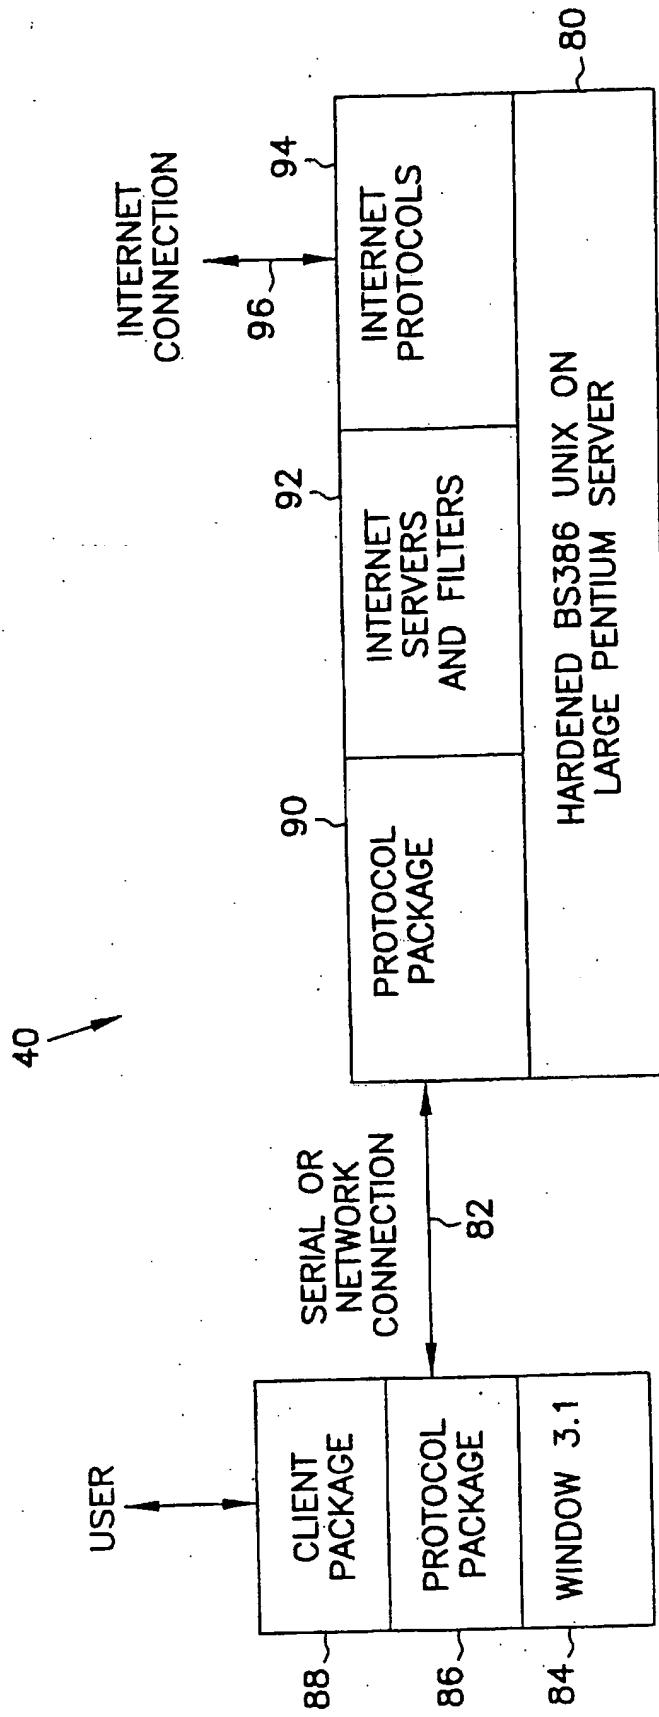


FIG. 4

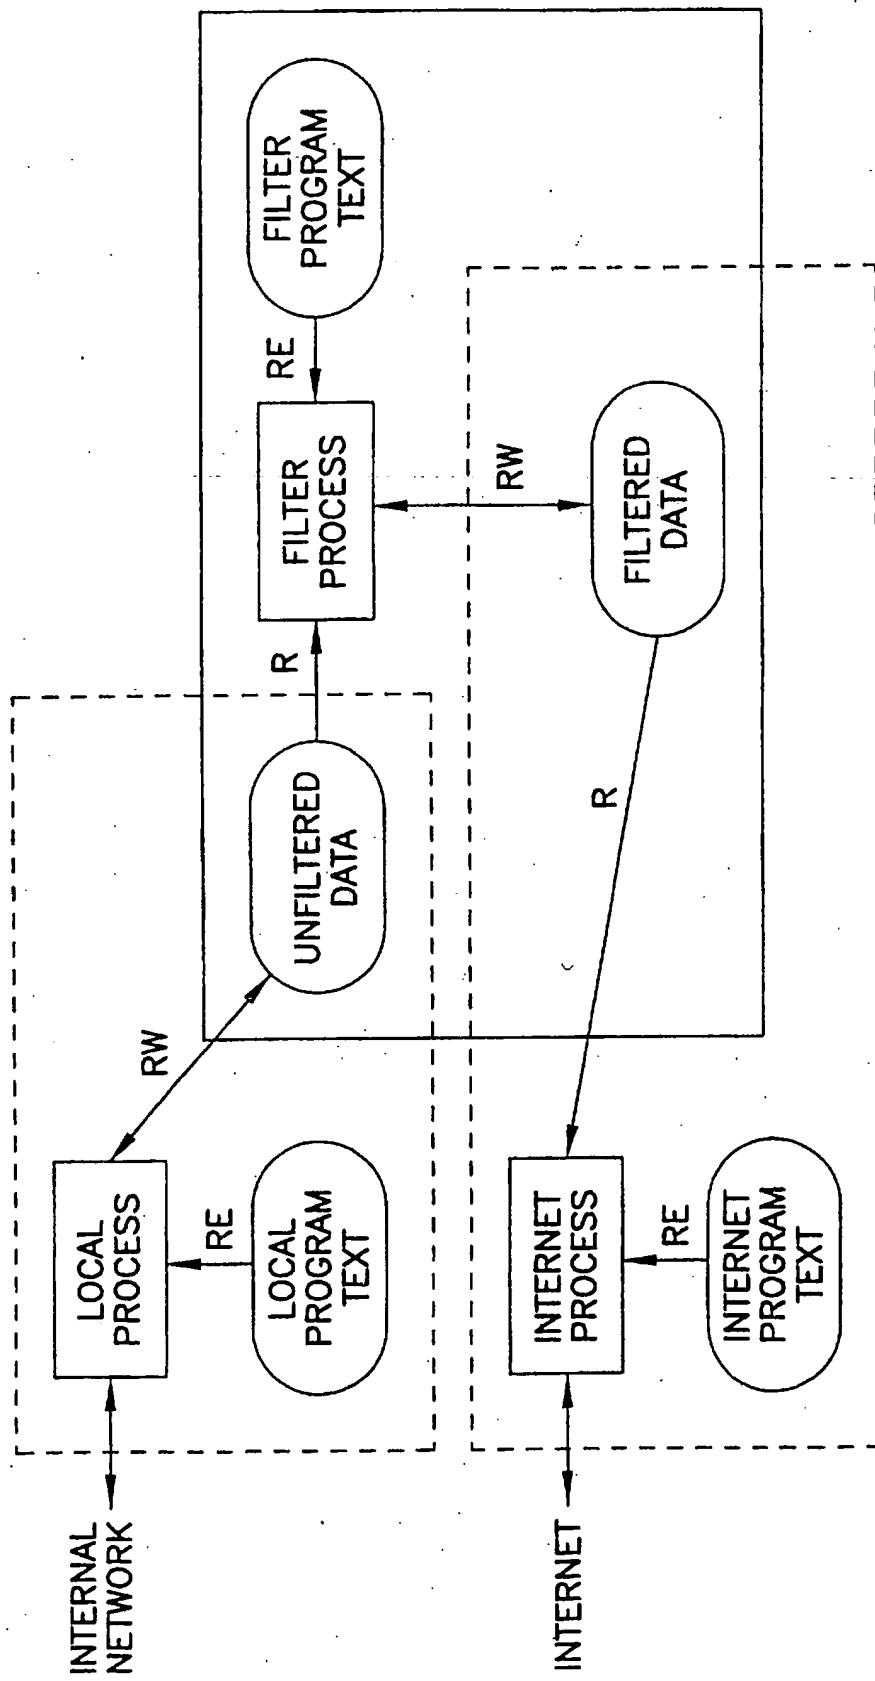


FIG. 5A

6 / 14

TYPE →

DOMAIN	UNFIL	DATA	FIL	DATA	CLIENT	PGM	FLTR	PGM	NET	PGM
LOCAL		RW				RE				
FILTER		R		RW				RE		
INTERNET					R				RE	

↓ →

R: READ ONLY  
RW: READ AND WRITE  
RE: READ AND EXECUTE  
████: NO ACCESS ALLOWED

FIG. 5B

7 / 14

Subtype	Source Code Name	Usage
file	file	Files that are private to the creating Domain.
directory	diry	Directories; not checked.
socket	sock	Reserved.
fifo	fifo	Fifos; not checked.
device	devi	Reserved.
port	port	Reserved.
executable	exec	Executable; effective Domain is not changed.
gate	gate	Executable; effective Domain is set to creator field of full Type name for duration of execution.

FIG. 6

Attribute Name	Meaning
ddt_read	Process may read (fetch) data from object.
ddt_write	Process may modify object.
ddt_rename	Process may rename object.
ddt_exec	Process may execute contents of object. Will only be assigned to subtypes gate and exec and will never be combined with ddt_write.
ddt_trigger	Trigger an alarm signal to Rover monitoring facilities as a side effect of granting access.
ddt_chcreator	If effective Domain of the process = creator field of Type, then process can change creator field.
ddt_destroy	Process may destroy the object.

FIG. 7

Syscall	Type of Executable Argument	Old Real Domain	Old Effective Domain	New Real Domain	New Effective Domain	Remarks
fork	No Argument	Mail	Mail	Mail	Mail	Child Process Spawed
execve	Mail:exec	Mail	Mail	Mail	Mail	New Executable
execve	\$Sys:exec	Mail	Mail	Mail	Mail	New Executable
make domain in with MIME as domain name	Mail:exec	Mail	Any	MIME	MIME	New Executable, as with execve
make domain with MIME as domain name	SMTP:gate	Mail	Any	MIME	SMTP	New Executable, as with execve
change domain with MIME as domain name	No Argument	Mail	Any	MIME	MIME	Continue With Same Executable
execve	SMTP:gate	Mail	Any	Mail	SMTP	Implicit Gating, New Executable
gate with SMTP as domain name	No Argument	Mail	Any	Mail	SMTP	Explicit Gating, Continue With Same Executable
ungate	No Argument	Mail	SMTP	Mail	Mail	Explicit Exit from Gate

FIG. 8

10 / 14

Privilege Name <sup>a</sup>	Meaning
can_ch_type	Can execute the ch_type or fch_type syscalls.
suppress_su_alarm	Can exercise super-user privilege without tripping alarm.
admin_reboot	Can transition to admin ("Greendome") mode.
can_set_clock	Can set or adjust the system clock.
can_setlogin	Can set user login name.
is_startup	Can perform functions required for startup only.

a. Note that there is no "can\_ch\_domain" privilege: the ability to change either the real or effective Domain of a process is controlled by the DIT.

**FIG. 9**

11 / 14

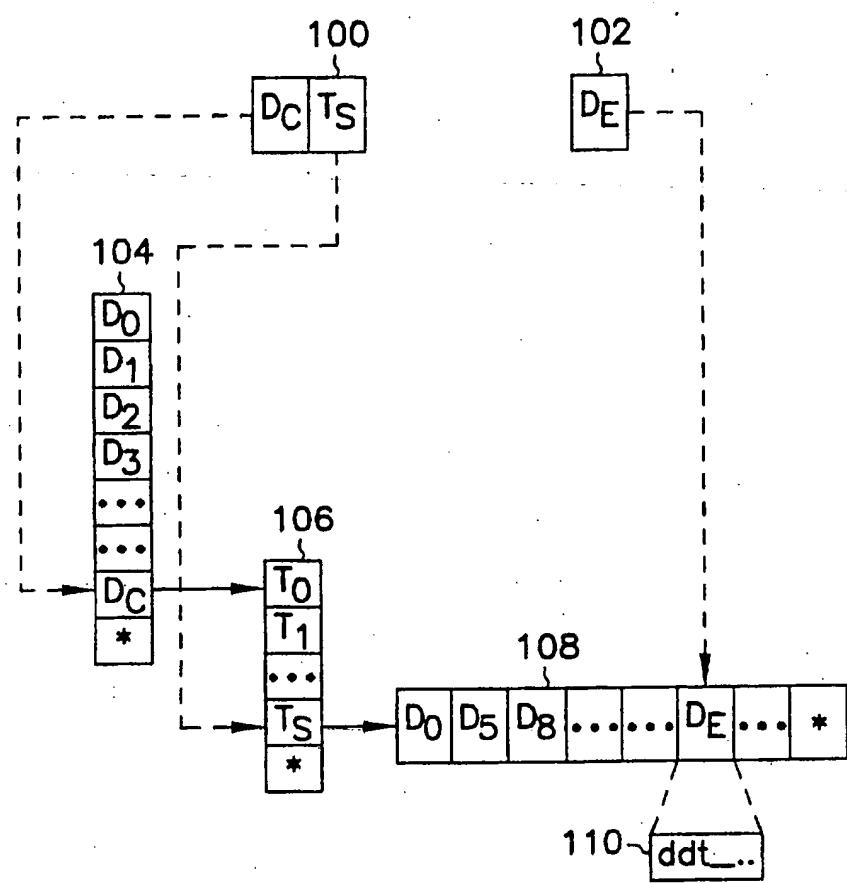


FIG. 10

12 / 14

D <sub>0</sub>
D <sub>1</sub>
D <sub>2</sub>
D <sub>3</sub>
...
...
D <sub>C</sub>
*

LIST OF ALLOWED D<sub>E</sub>  
THAT D<sub>1</sub> CAN GATE INTO

FIG. 11

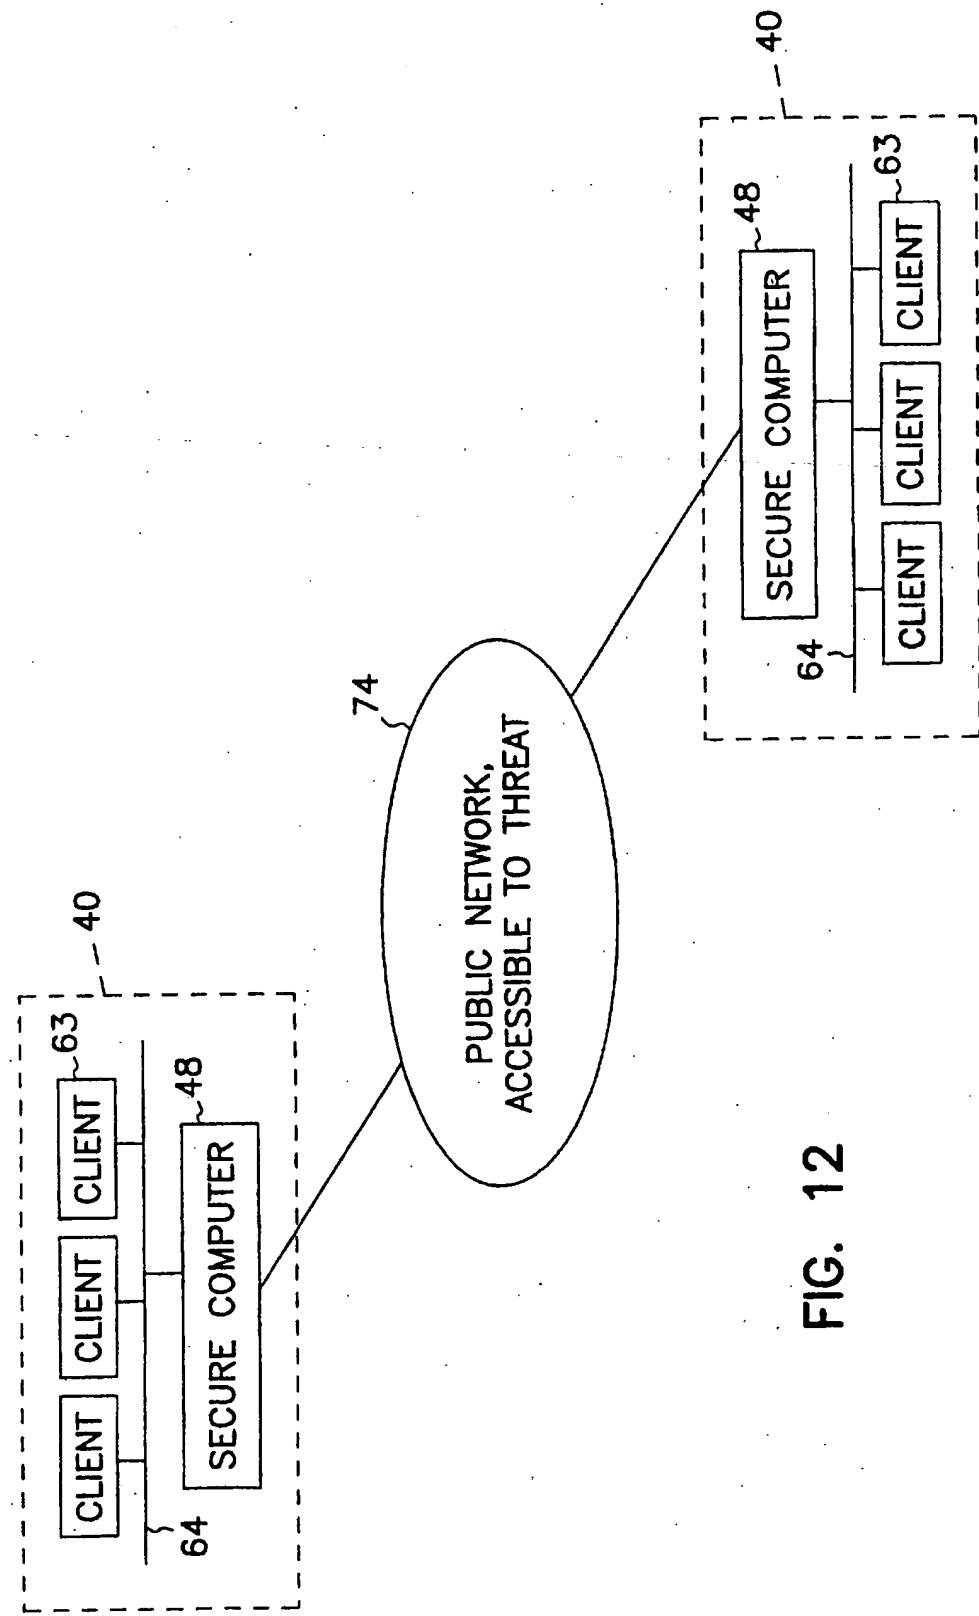


FIG. 12

14 / 14

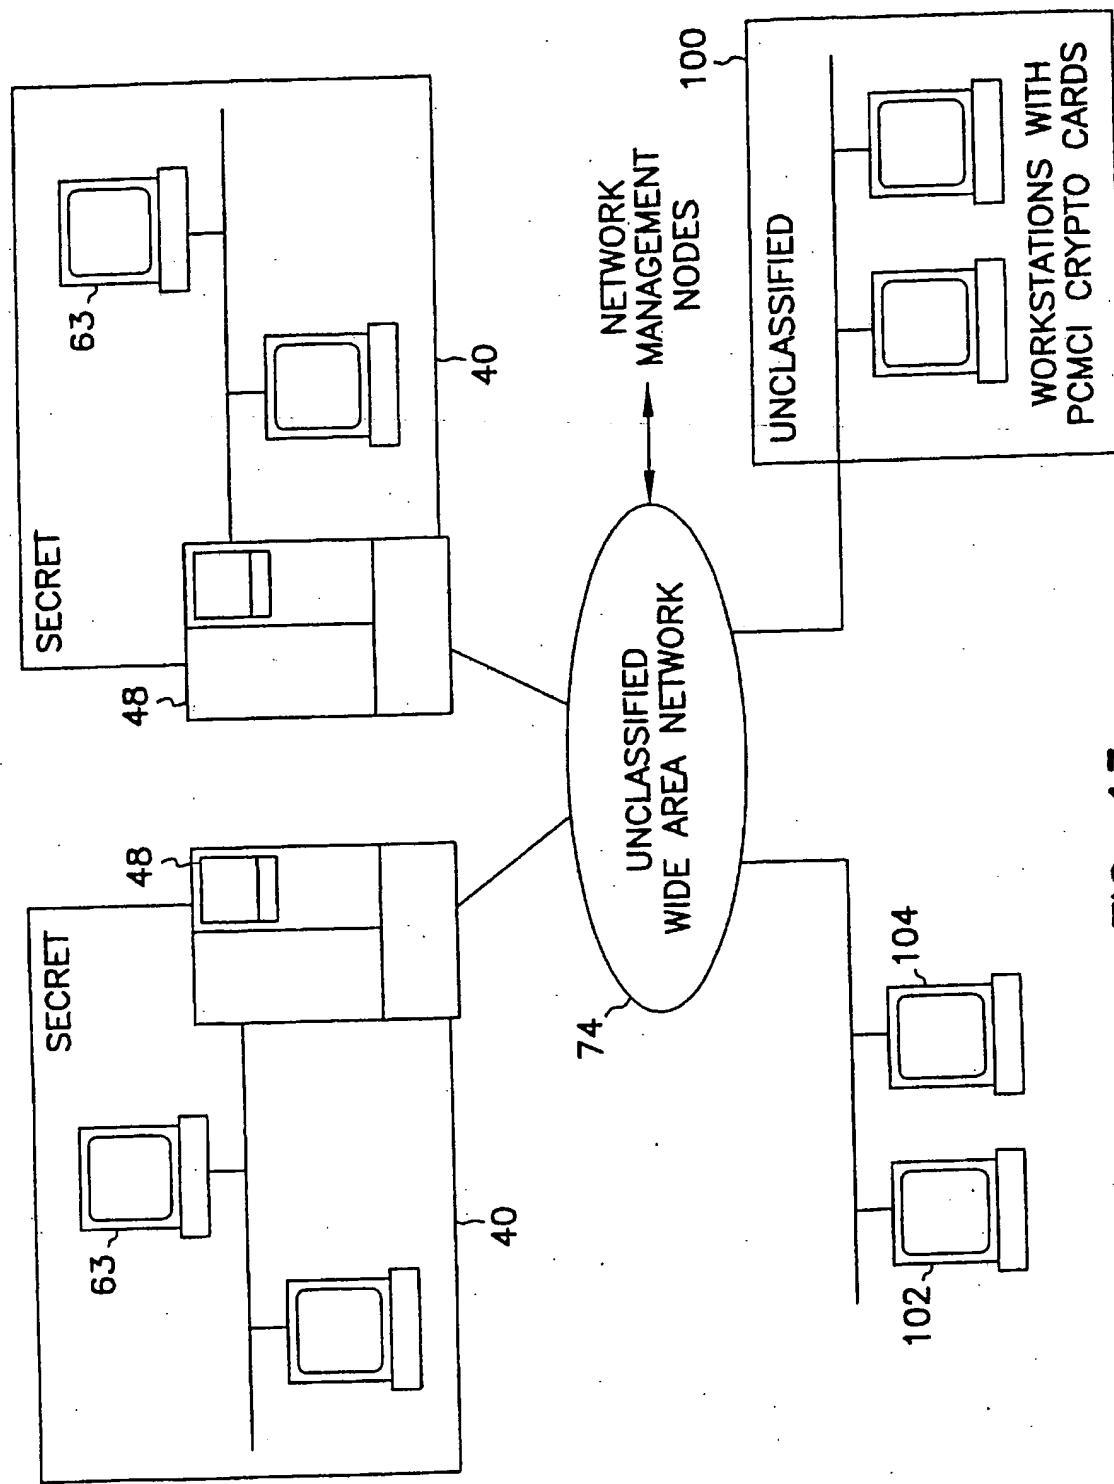


FIG. 13

**This Page is Inserted by IFW Indexing and Scanning  
Operations and is not part of the Official Record**

## **BEST AVAILABLE IMAGES**

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADED TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER:** \_\_\_\_\_

**IMAGES ARE BEST AVAILABLE COPY.**

**As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.**